Fermilab

Managed by Fermi Research Alliance, LLC for the U.S. Department of Energy Office of Science

HEP Facility Plans

Lothar A. T. Bauerdick
U.S. CMS Operations Program Manager

June 10, 2015

Distributed Facilities and Sharing of Resources

- Experimental HEP mostly uses **distributed high-throughput computing**
  - ✦ Before LHC most of the computing and storage capacity was at a central experiment data center
  - ✦ now most HEP experiments have adopted distributed computing, with global workflow, scheduling, and data management enabled by high-performant networks and a distributed operations and security infrastructure

WLCG

Open Science Grid

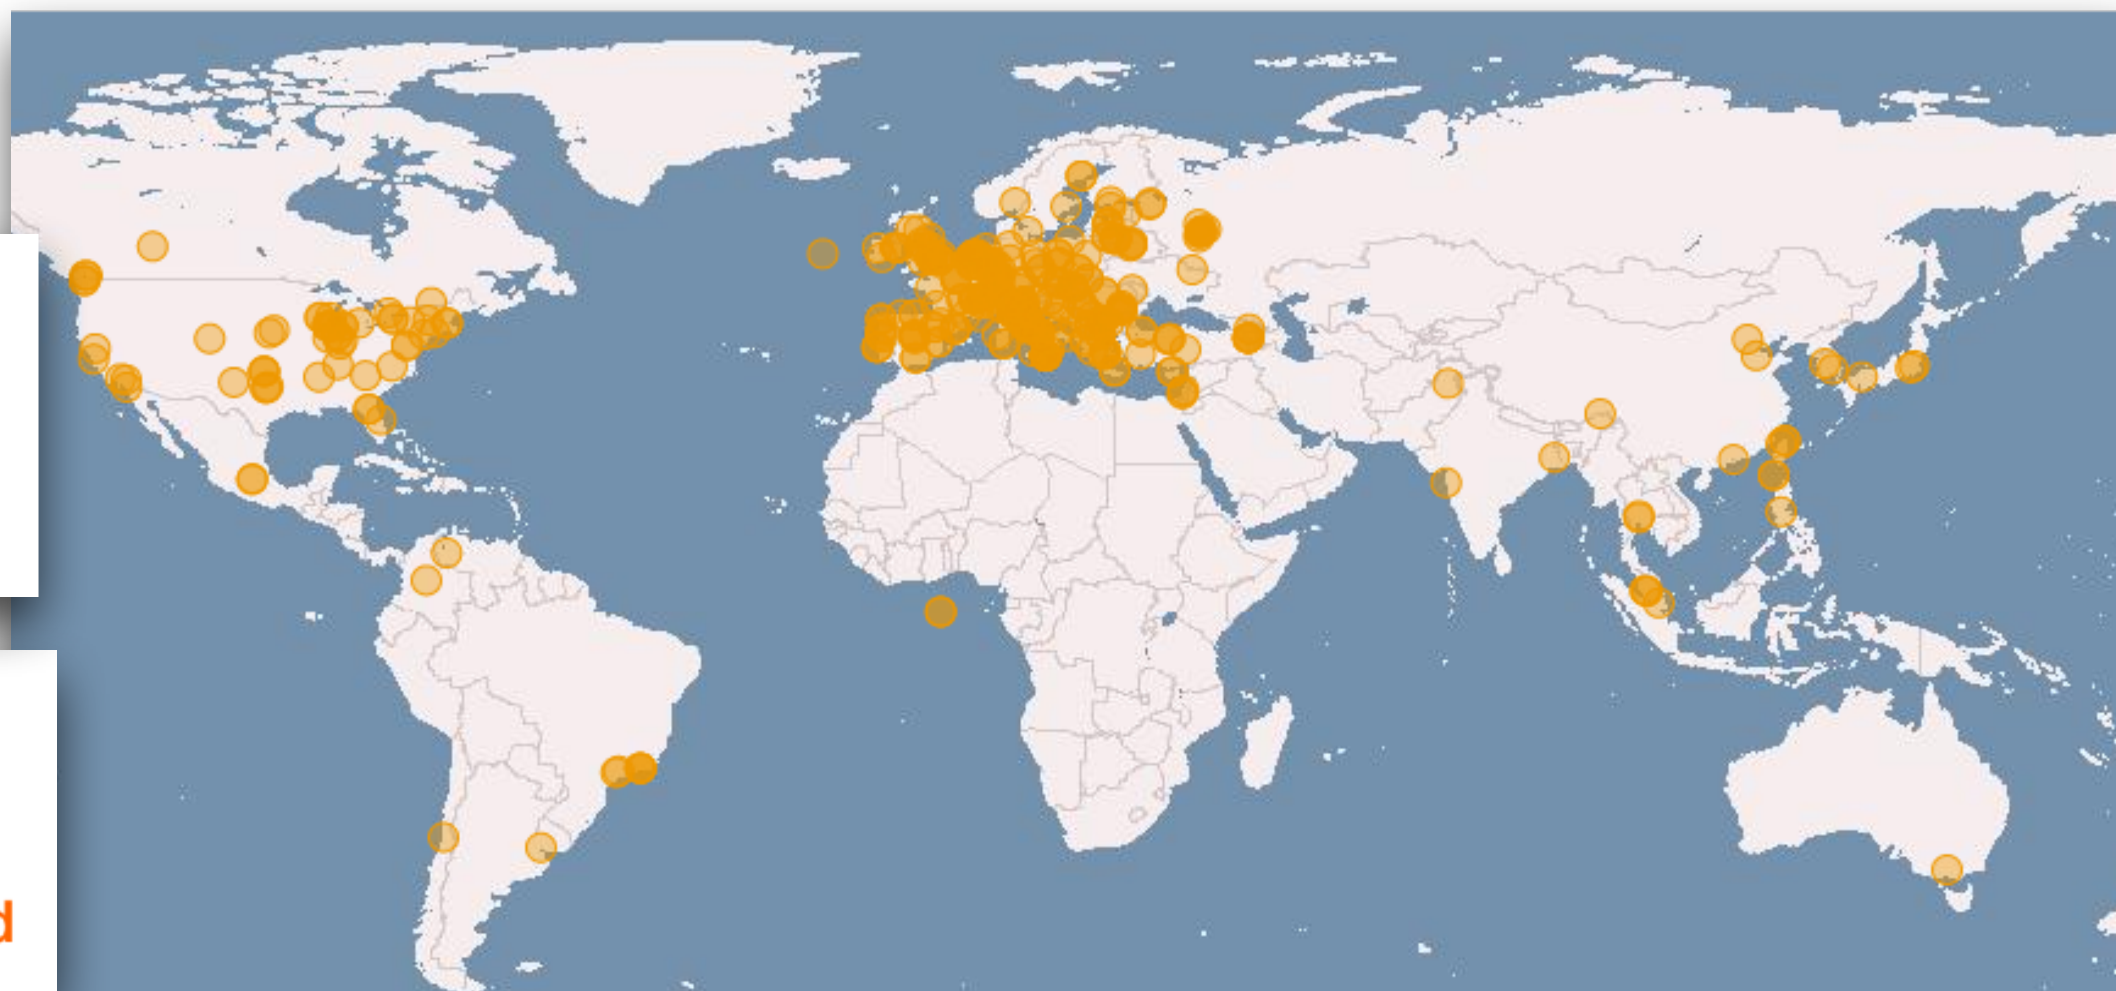

Facilities Services to Support HEP

- **High Throughput Computing** facilities: batch system across worker nodes
  - ✦ **HTC** means shared utilization of large ensembles of autonomous resources, where all elements are optimized for maximizing throughput
  - ✦ HEP tasks (data analysis, reconstruction, simulation) split into parallel jobs
    - no/loose inter-process communications, it's an "embarrassingly parallel" problem
  - ✦ distributed computing resources, presented to users/work flow engines through "overlays" as a coherent job execution environment
    - pilot-job based Overlay Job Managers through glideinWMS, Panda, HTcondor
- **Work Flow** aspects are very important for the quality of the scientific result
  - ✦ e.g. significant data management, book keeping, provenance tracking
  - ✦ requires facilities to provide services to support these capabilities
  - ✦ including schedulers and workflow systems, information systems, data bases, trust relationships and security protocols, distributed monitoring and problem resolution
- Very significant **Data Management** and **Data Access** component
  - ✦ managing and extracting science from tens of PB active data, created and hosted elsewhere
    - simulations data observational ~ same size or larger and more diverse
    - expect to grow to ~.3 ExaByte of active data in the US during this decade
  - ✦ support federation of storage systems, high performance networks, data management management systems, across a set of data centers for local and remote access to data

HEP Scientific Facilities at US Labs and Universities: A rich stack of distributed services on top of “bare metal”

- Provide and manage computing services and resources
  - Data recording, storage, access, bulk processing, analysis
  - CPU Cores, Online (Disk) and Offline (Tape) Storage, Networking
- HEP Labs and HPC support specific and often wide-range distributed communities
  - FNAL Tier-1 facilities for LHC/CMS and for Intensity Frontier experiments, BNL Tier-1 for Atlas and for RHIC experiments, PNNL is US facility for Belle, NERSC and other HPCs for large simulation etc
  - The program also supports large number of university computing Tier-2 and Tier-3 centers for LHC
- Large needs require large capacity centers (example: resources at Fermilab)

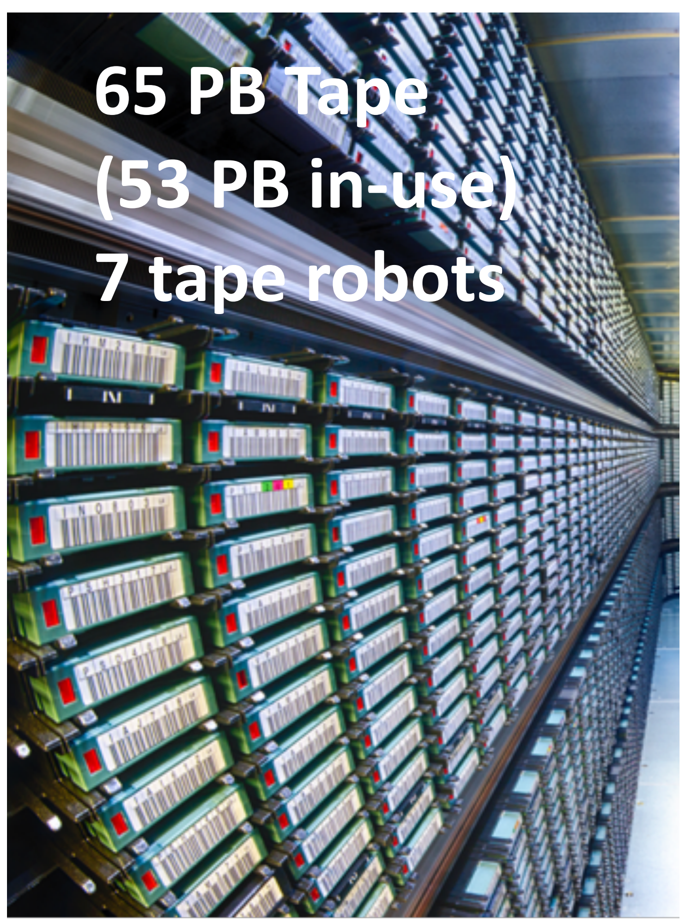

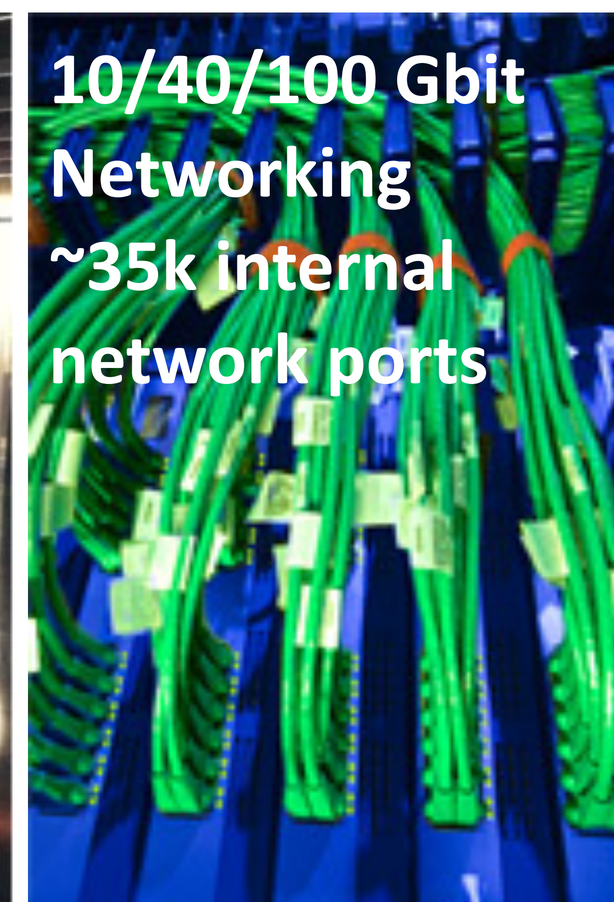

Example: Scientific Computing Facilities at Fermilab

High Throughput Facility: lightly-coupled workflows

Scientific Data Management

Scientific Data Storage and Access

CMS Facility

General Purpose Facility

Access to OSG

High Performance Facility: strongly-coupled workflows

US Lattice QCD, Accelerator Modeling, Cosmology

Virtual Facility

Provisioning static and dynamic (on demand) resources via the same interfaces

Extending facility footprint outside of FNAL (clouds, grid) transparently to the users

R&D

Evaluating emerging technologies for suitability for scientific workflows

e.g.:
GPGPUs
CPU accelerators
Object stores

Developing tools & costing model for optimal utilization of clouds (commercial, …) vs local resources

Example: ATLAS Distributed Computing

from M.Ernst

US Tier-2s
AGLT2: U Michigan, Michigan State
MWT2: UoChicago, Indiana U, UIUC
NET2: Boston U, Harvard
SWT2: UTA, Oklahoma U
WT2: SLAC

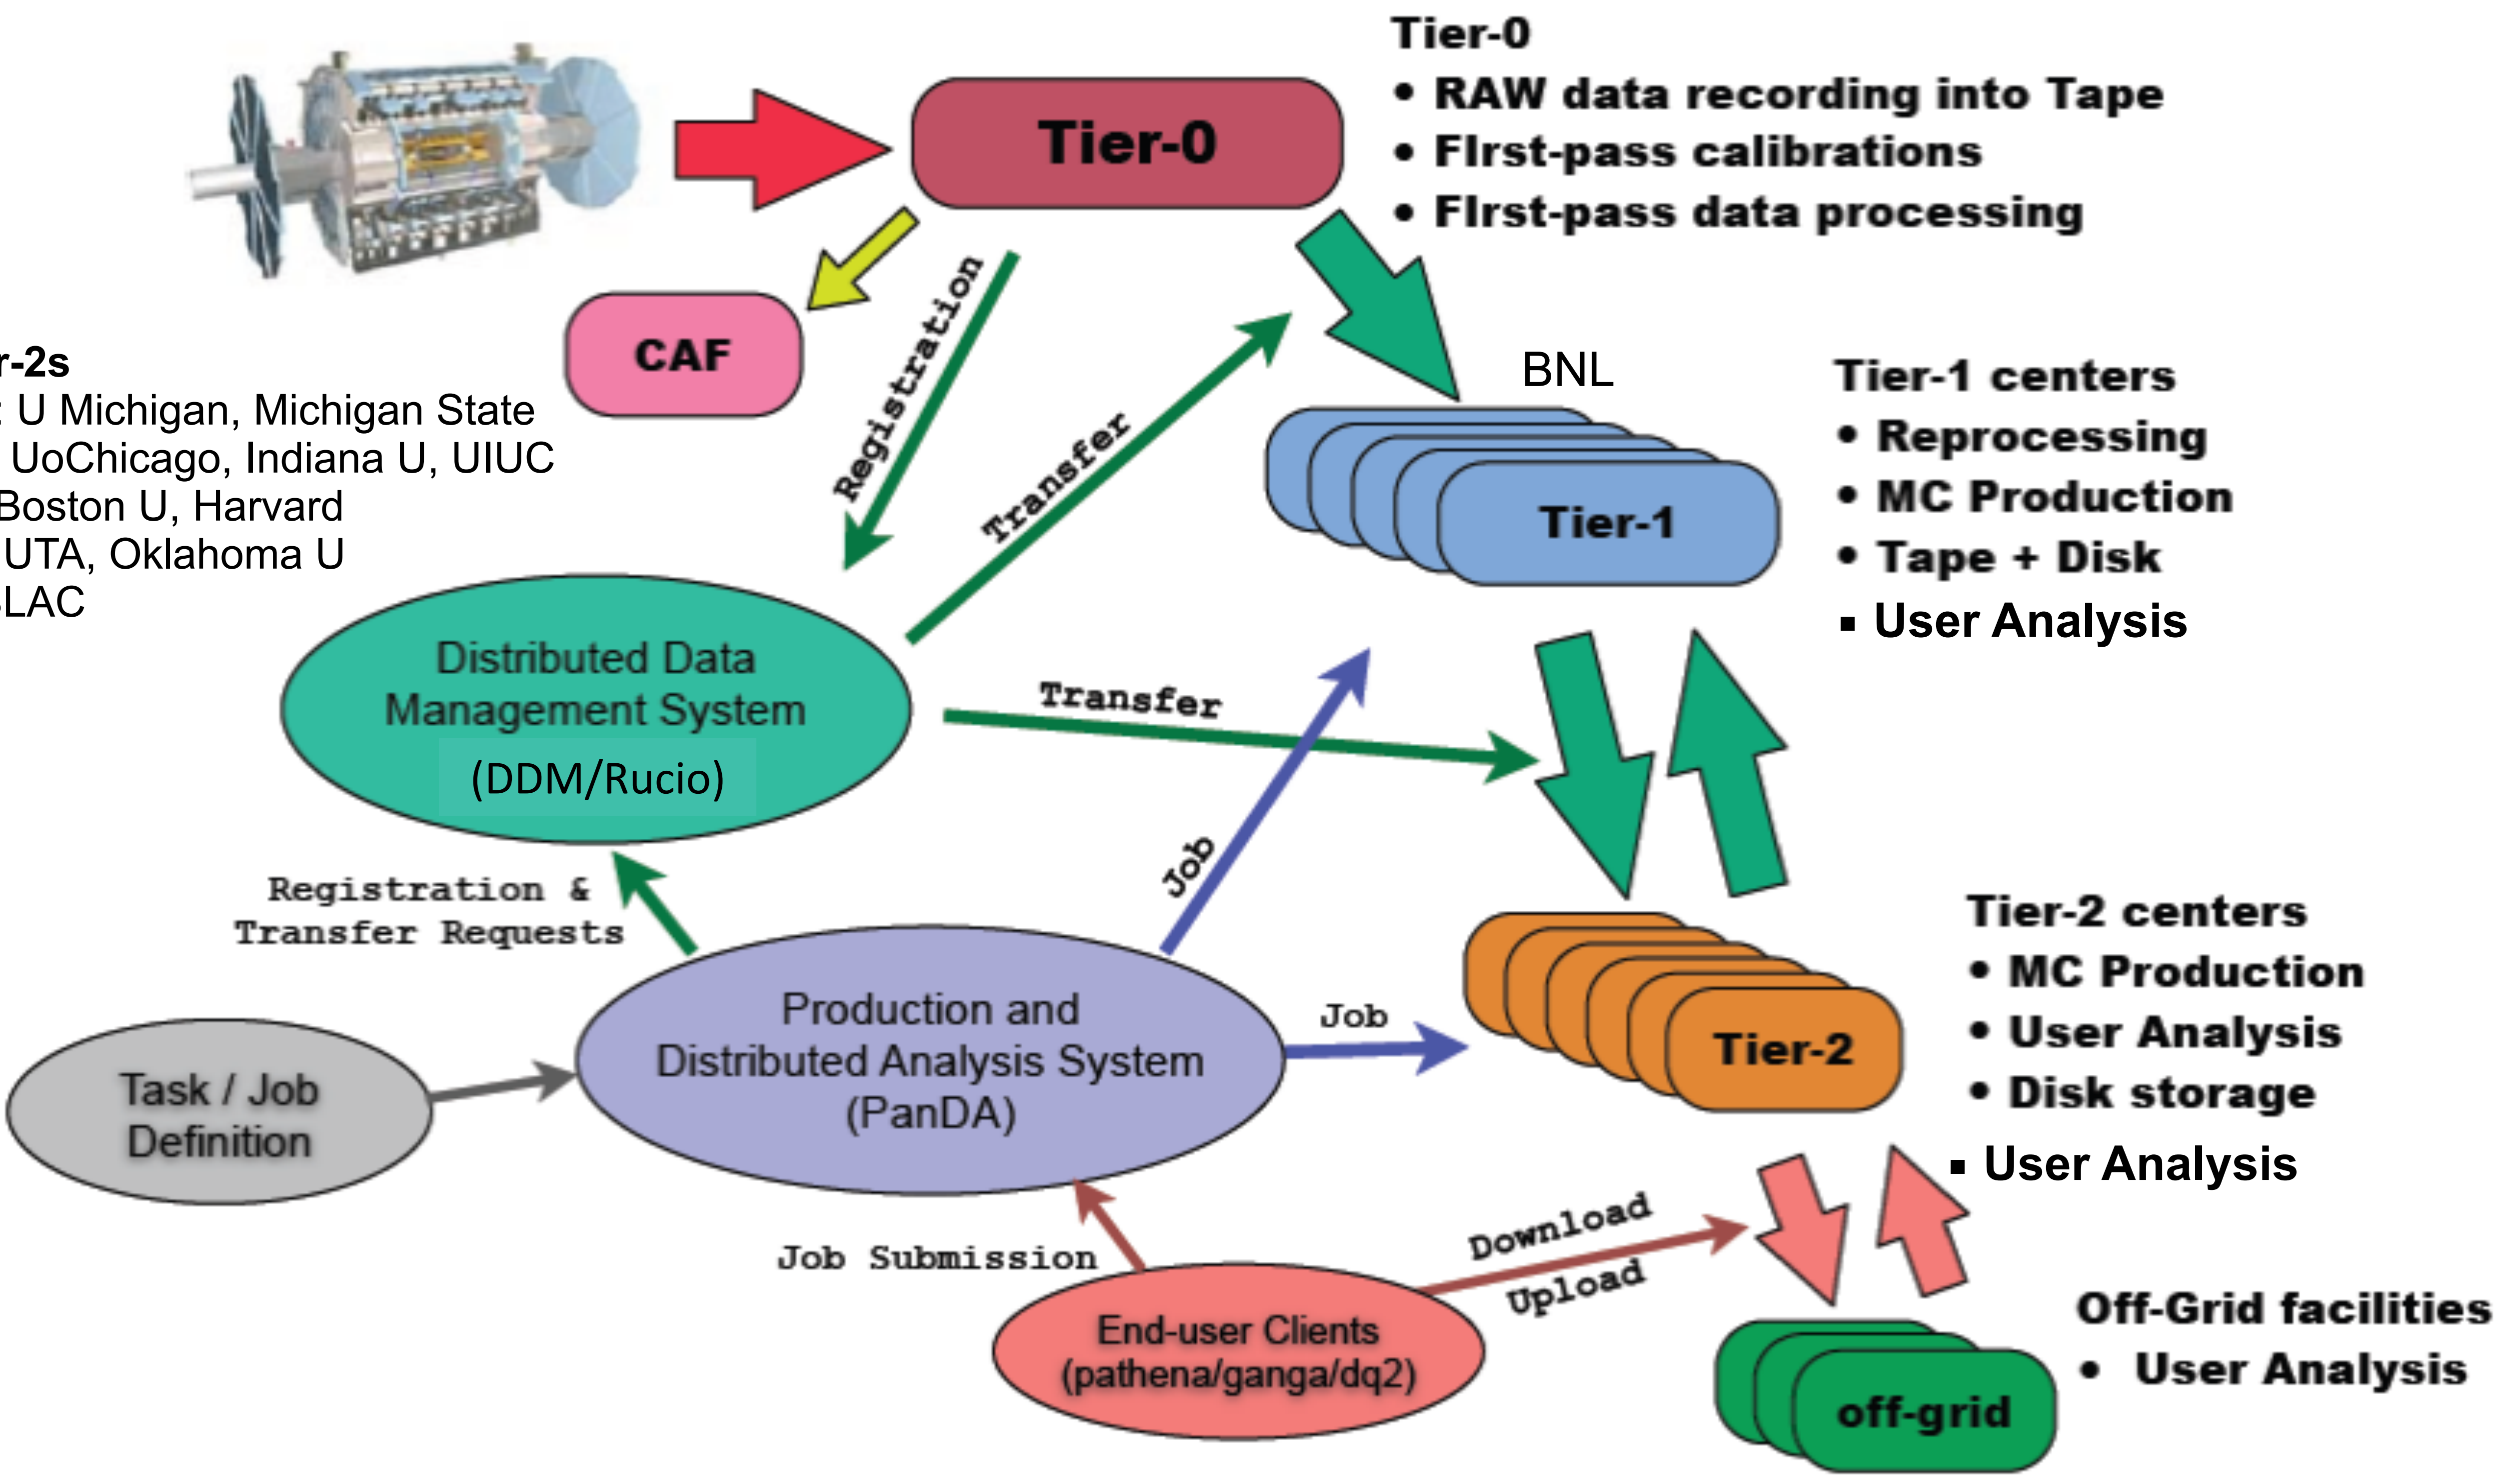

Significant shift toward a truly distributed environment

- Computing resources are intended to work more as a coherent system than a collection of sites with specialized functions
  - Improvements in data access across the Wide-Area (AAA, REX) and improved networks have been key to this!
- Sites present their capabilities to the distributed environment
  - as a set of computing services with certain performance characteristics (like #cores/node, memory/core, time limits, availability of special hardware, …)
  - as disk-based storage systems with certain data sets and capabilities
  - as long-term lower-cost archival systems with tape library backends
- Workflows and data management systems of science customers (VOs) work across this environment
  - The Tier-1 centers, OSG Grid Operations Center and other sites provide the glue and the service layer, including cyber security, user support and problem resolution etc.
- Thus, capabilities and capacities can be added when available
  - including new kind of resources like commercial clouds or allocation-based resources at HPC!

Facility Capacity Needs for Data Management, Data Storage and Data Access Systems

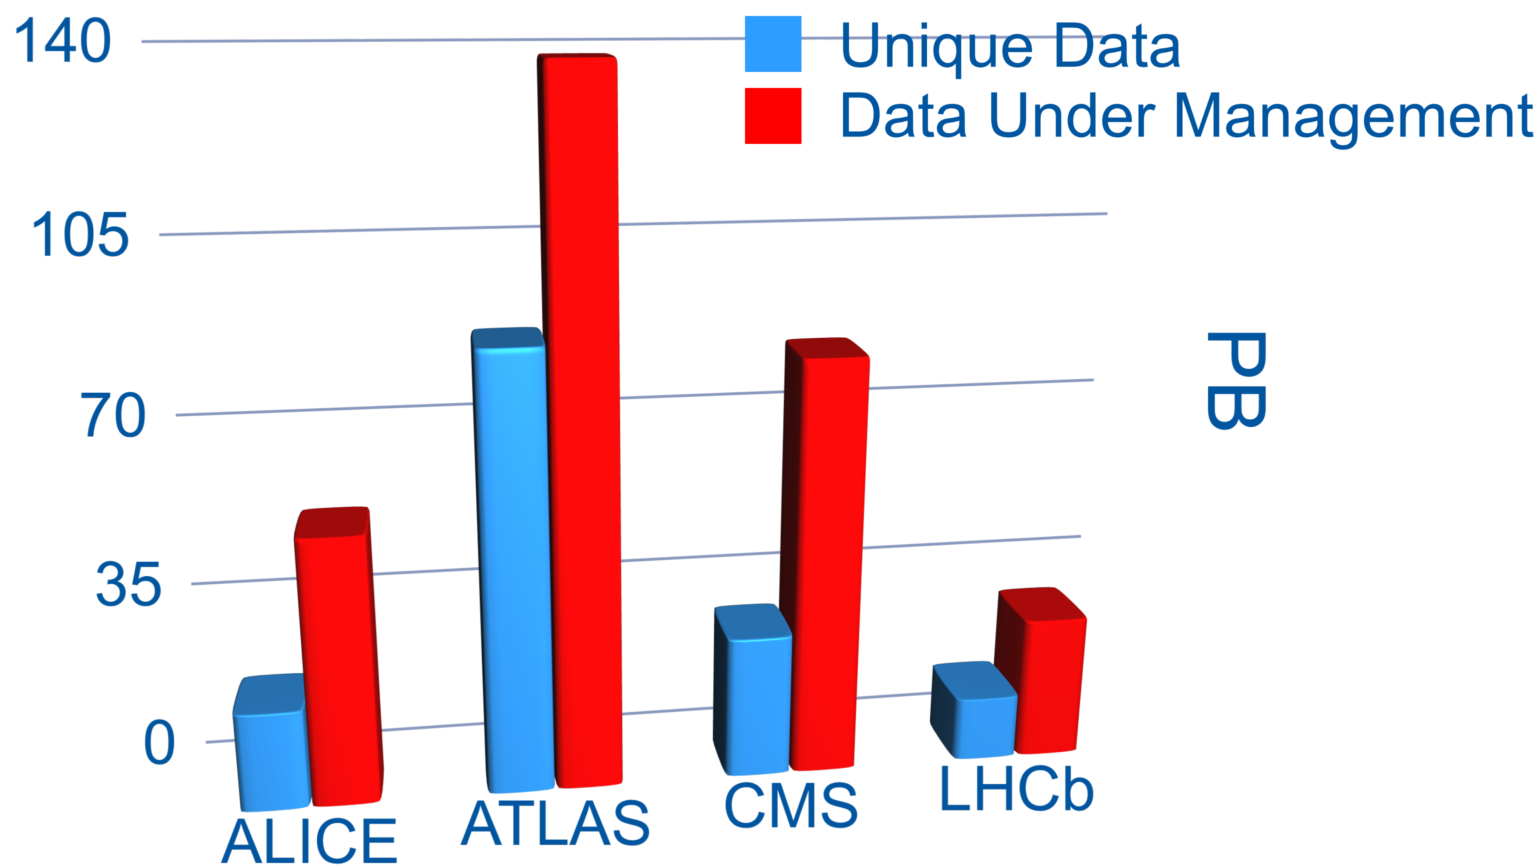

- There are close to 200 facilities contributing computing to LHC, hosting 140 PB of unique and a total of 280 PB managed data
  - data set dimensions: 246 PB of disk and 267 PB of tape
  - More than 1B files with average file size between 0.2GB and 2.5GB

Long history of provisioning and operating storage for all Fermilab Users: approaching 100 Petabytes of data…

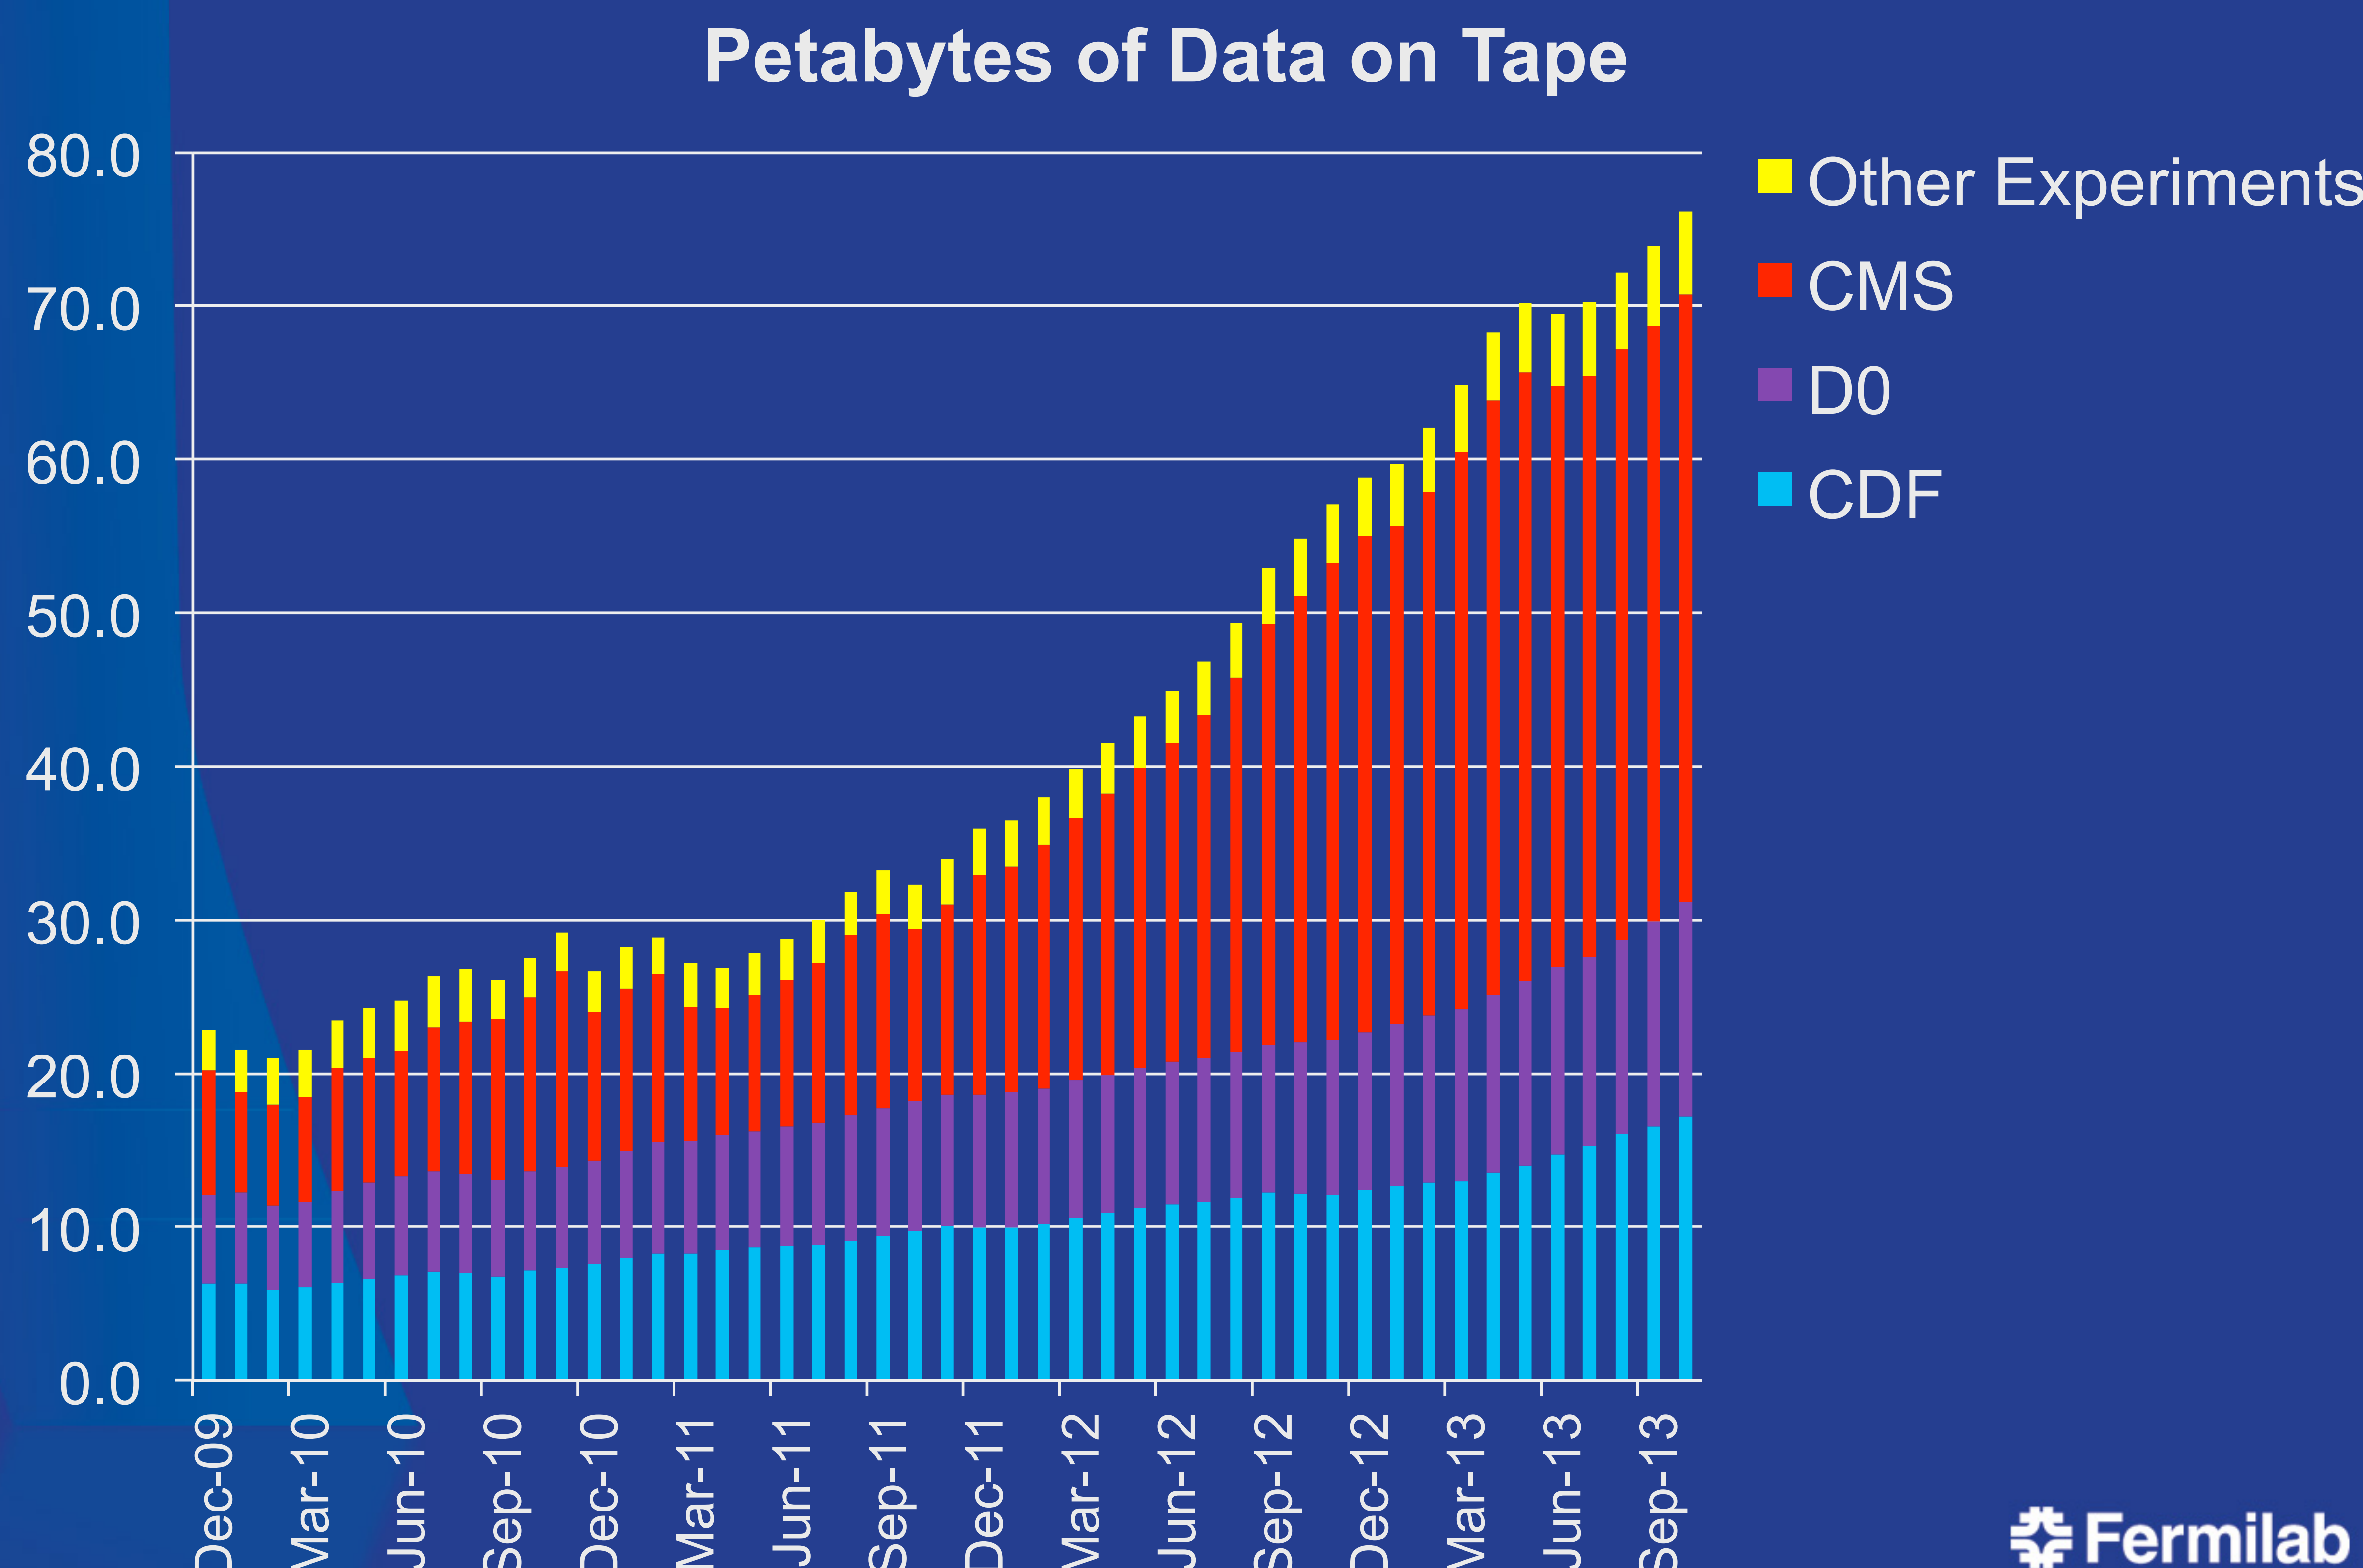

Fermilab

Facility Capacity Needs for HEP Computational Services
CMS CPU use for distributed simulation and data analysis

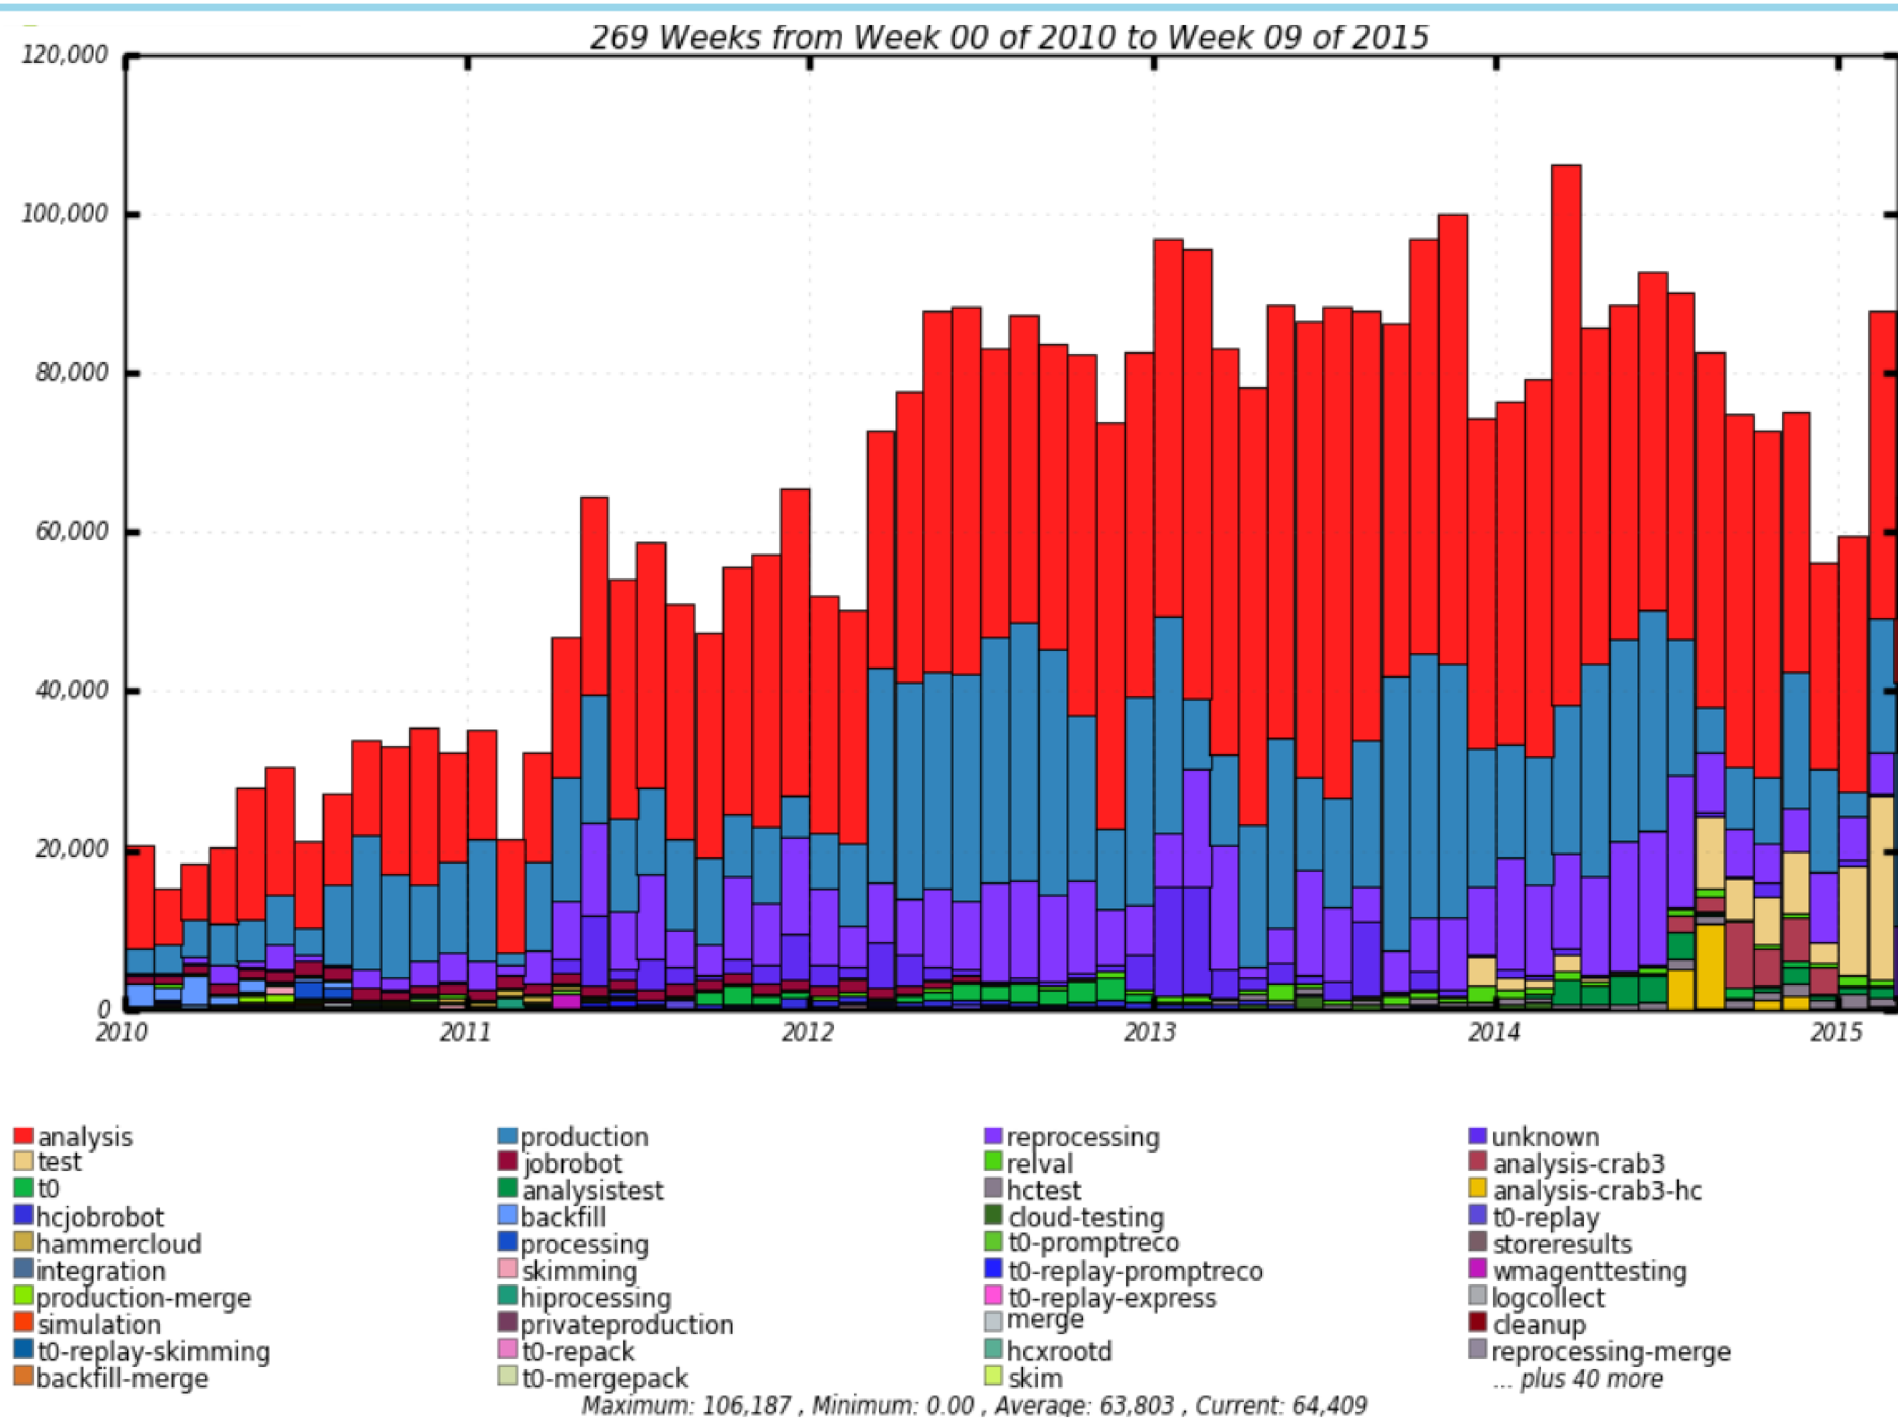

- CMS has ≈100k cores available worldwide that are useable for simulation production and analysis of datasets

Fermilab

Atlas Sustained Usage (World)

from M.Ernst

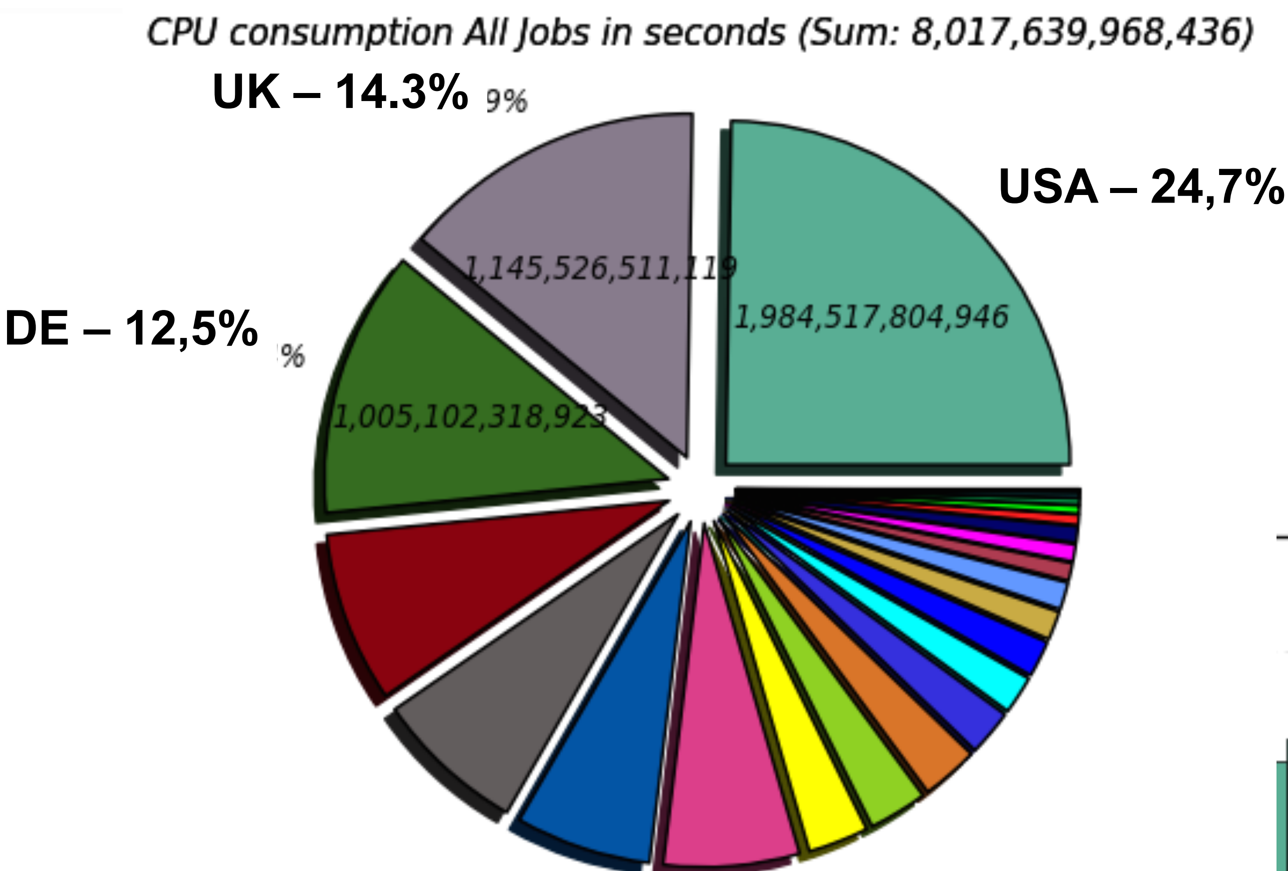

Atlas: 890 Million CPU hours per year over the LHC shutdown
220 Million hours / yr in the U.S.

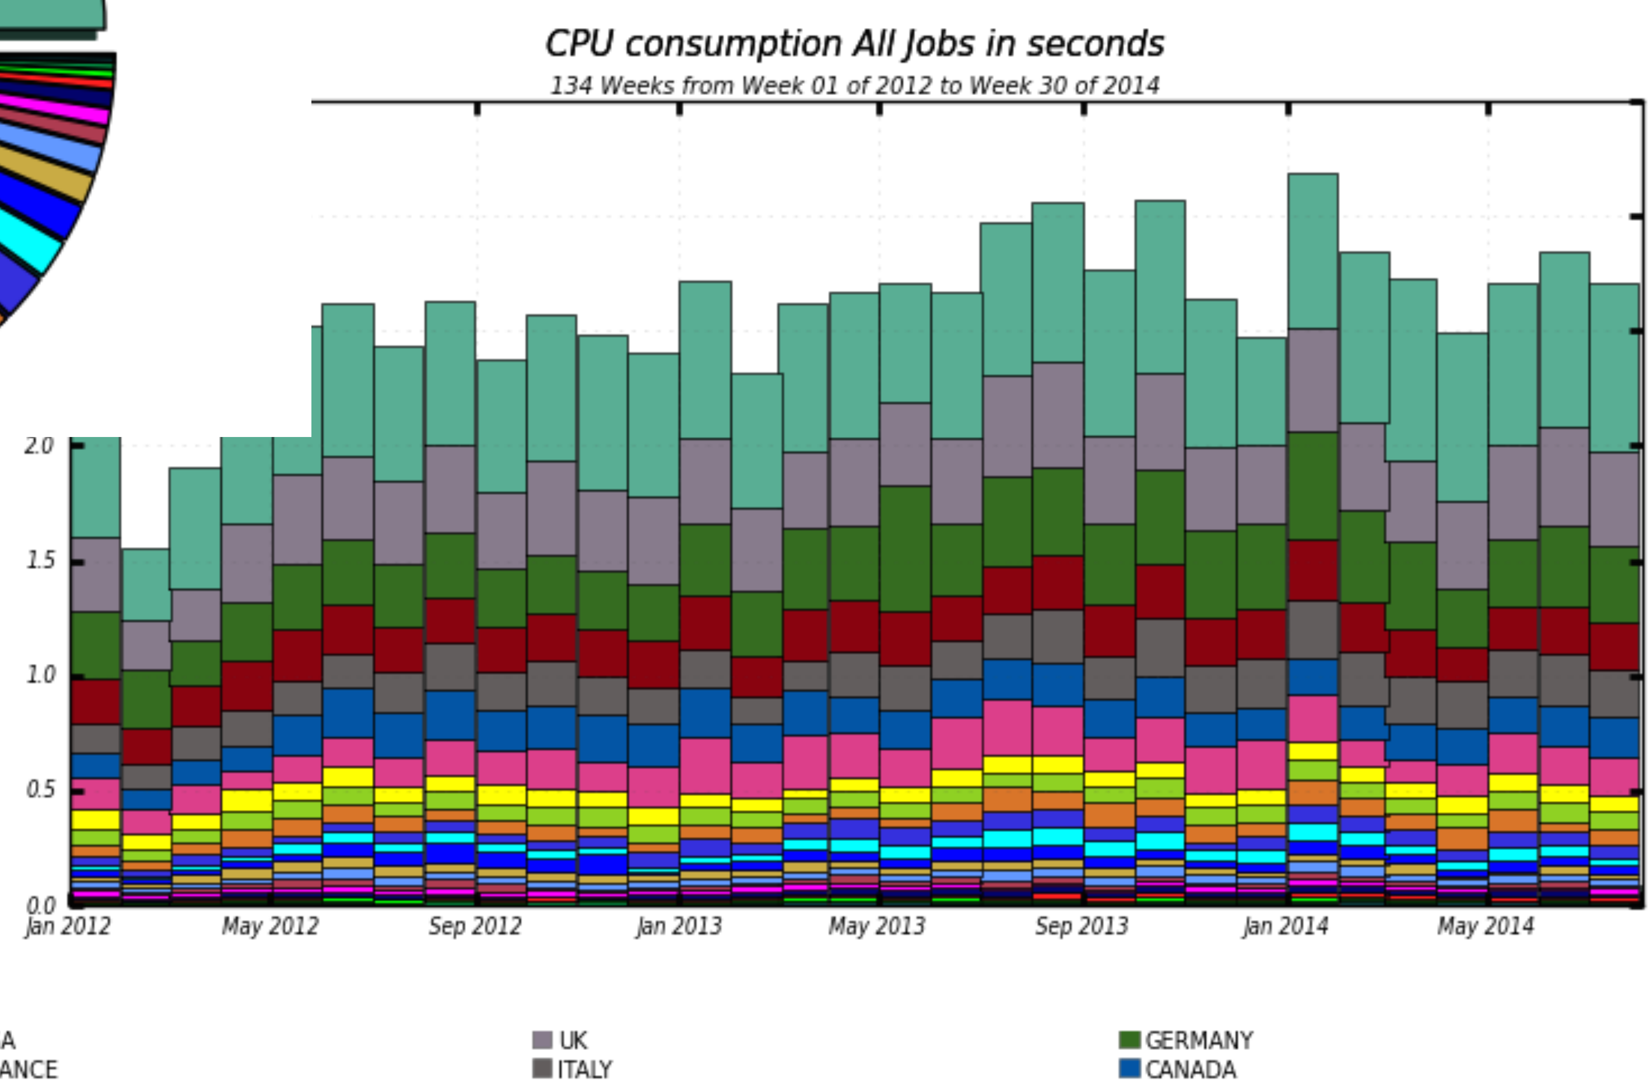

both Atlas and CMS about 100,000 cores
2-4 GB memory/core

US HEP Facilities are part of Open Science Grid

- OSG Delivers up to 3.5 Million CPU hours every day
  - about 60% go to LHC, 20% to other HEP, 20% to many other sciences
- OSG has a footprint on ~120 campuses and labs in the U.S.
  - Supports active community of 20+ multi-disciplinary research groups

Increasing size of US LHC computing and facilities

- U.S. LHC facilities are part of the Open Science Grid
- Since 2007, 3.8 Billion CPU hours delivered, more than a billion jobs run!
  - In past 12 months, 823 M CPU hours and 200 M Jobs
  - Provides access to O(150k) cores

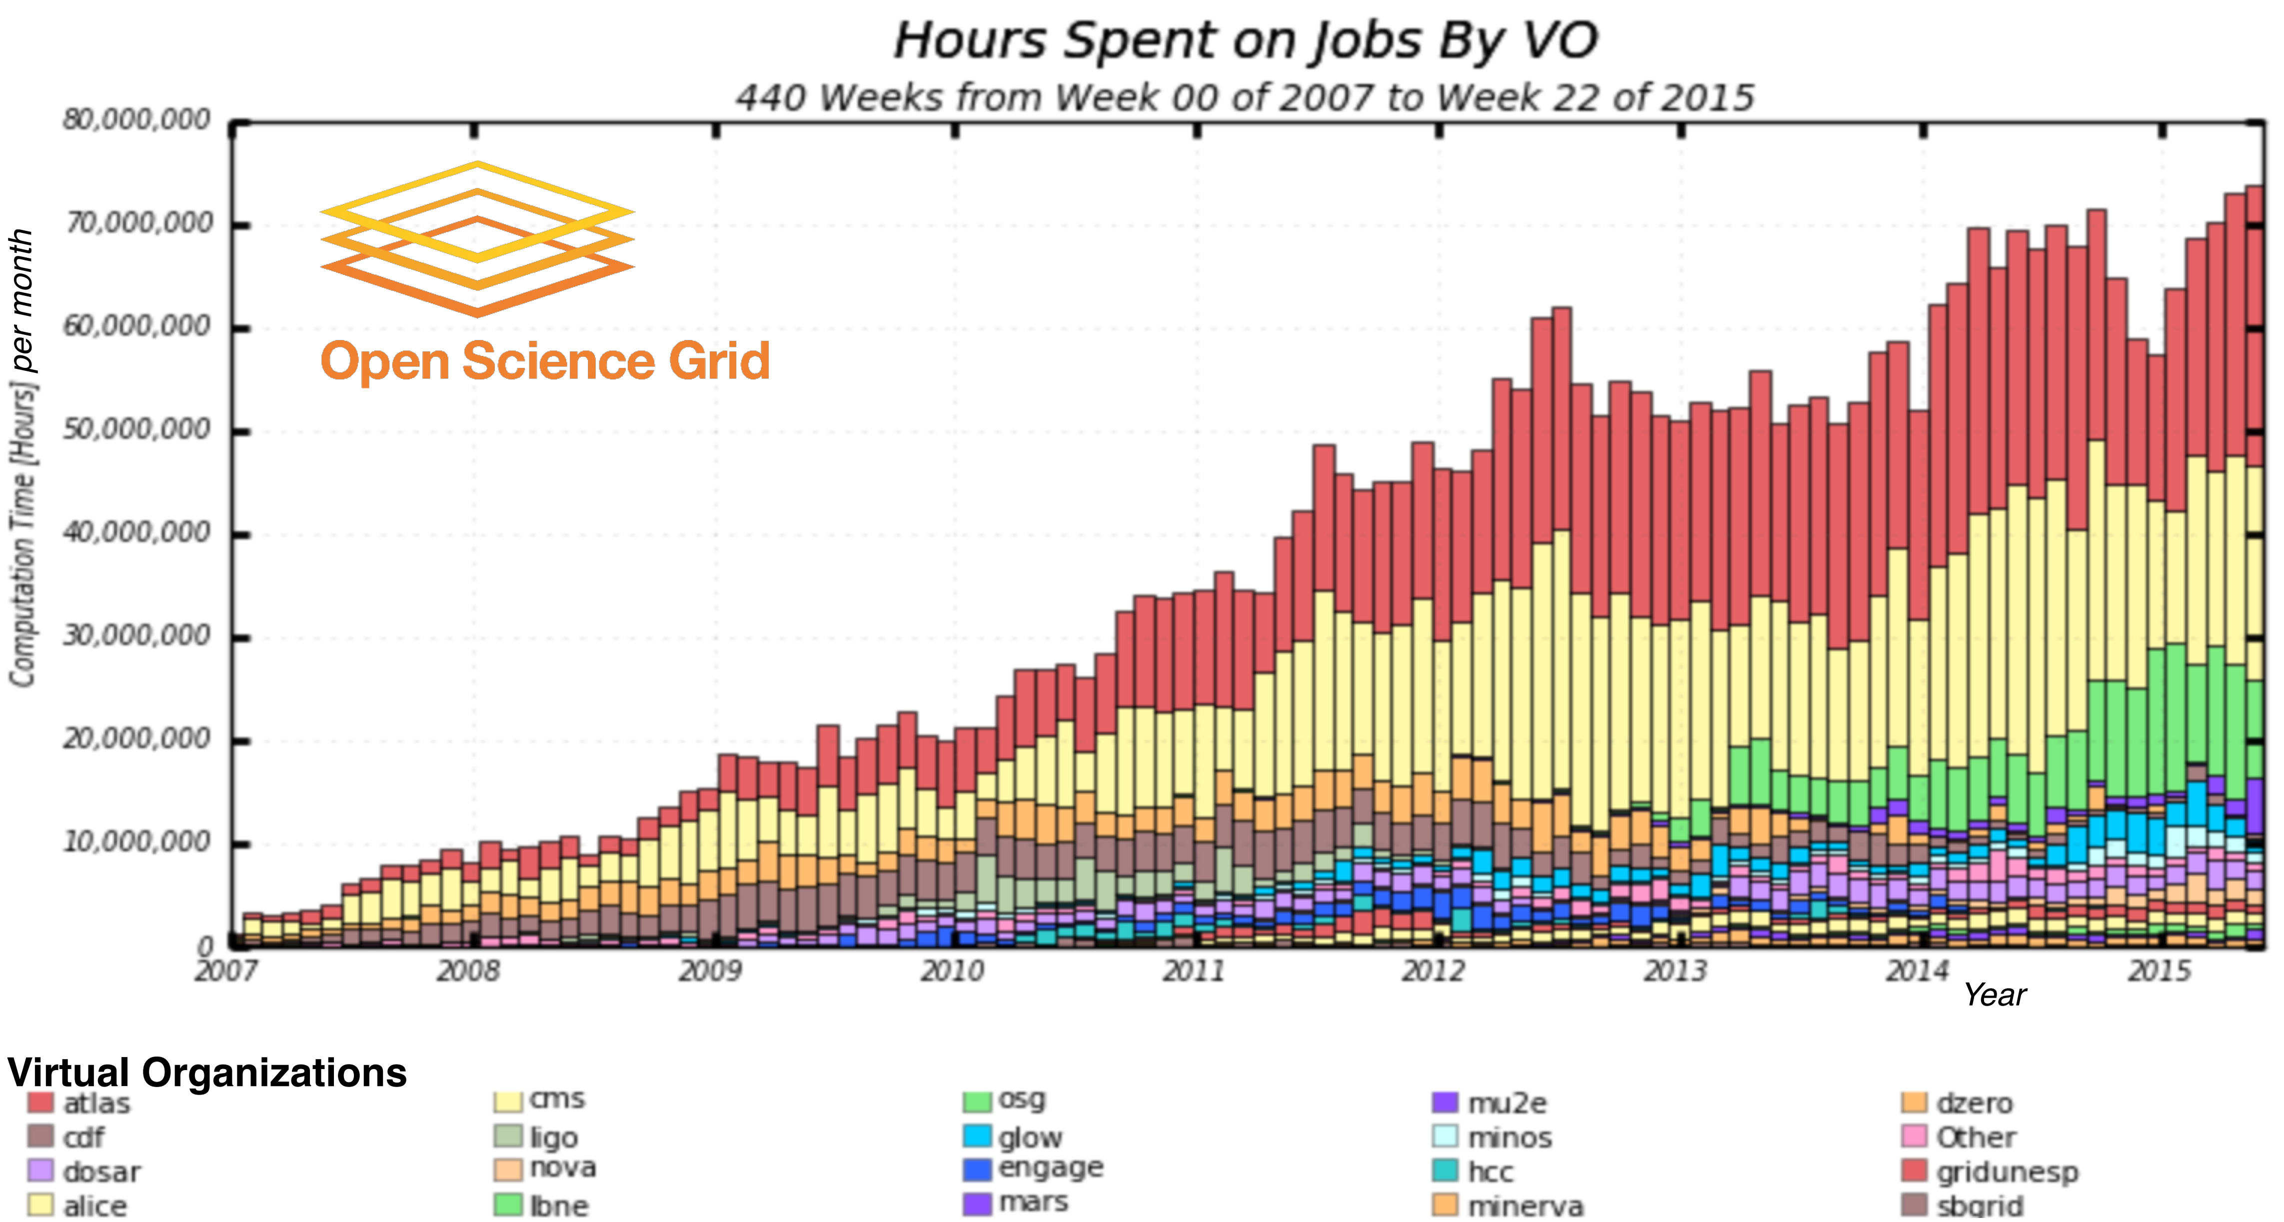

Distributed HEP Facilities rely on OSG for operations support

Open Science Grid

- **OSG Operations Group** (Indiana U., UCSD, Fermilab, and BNL)
  - ✦ provide the OSG platform/eco system of DHTC services, sites, software **to enable VOs** to run workflows and data systems **across OSG sites**
    - running a world-class unique diverse set of services, enabling more than 100 sites
    - infrastructure services, operations support, cyber security and incident response etc
    - main customer is the LHC, and other large experiments/VOs
  - ✦ provide a production quality **HTC facility** built on **harvesting resources opportunistically** from OSG sites, for a large & diverse community of researchers and science platforms
    - including significant users from DOE sciences
  - ✦ OSG provides “non-intrusive” ways of **connecting facilities** to the grid
    - important to science users at universities, and for access to non-HEP facilities
  - ✦ provide **other added value**:
    - user and host certificates (OSG CA, the follow-up of DOEgrids CA)
    - software distribution services (OASIS based on CVMFS)
    - network monitoring and dashboard
    - …

The OSG Harvesting Free CPU Cycles :-)
>100M hours provided this way, past 12 months

Open Science Grid

Fermilab

HEP Computing Resource Planning: Detailed Modeling of managed workflows and average analysis use

- Example CMS:
  - Very large increase in online-storage (disk systems)
    - at Tier-1 (lab) and Tier-2 centers
  - This is a main cost driver for LHC computing upgrades

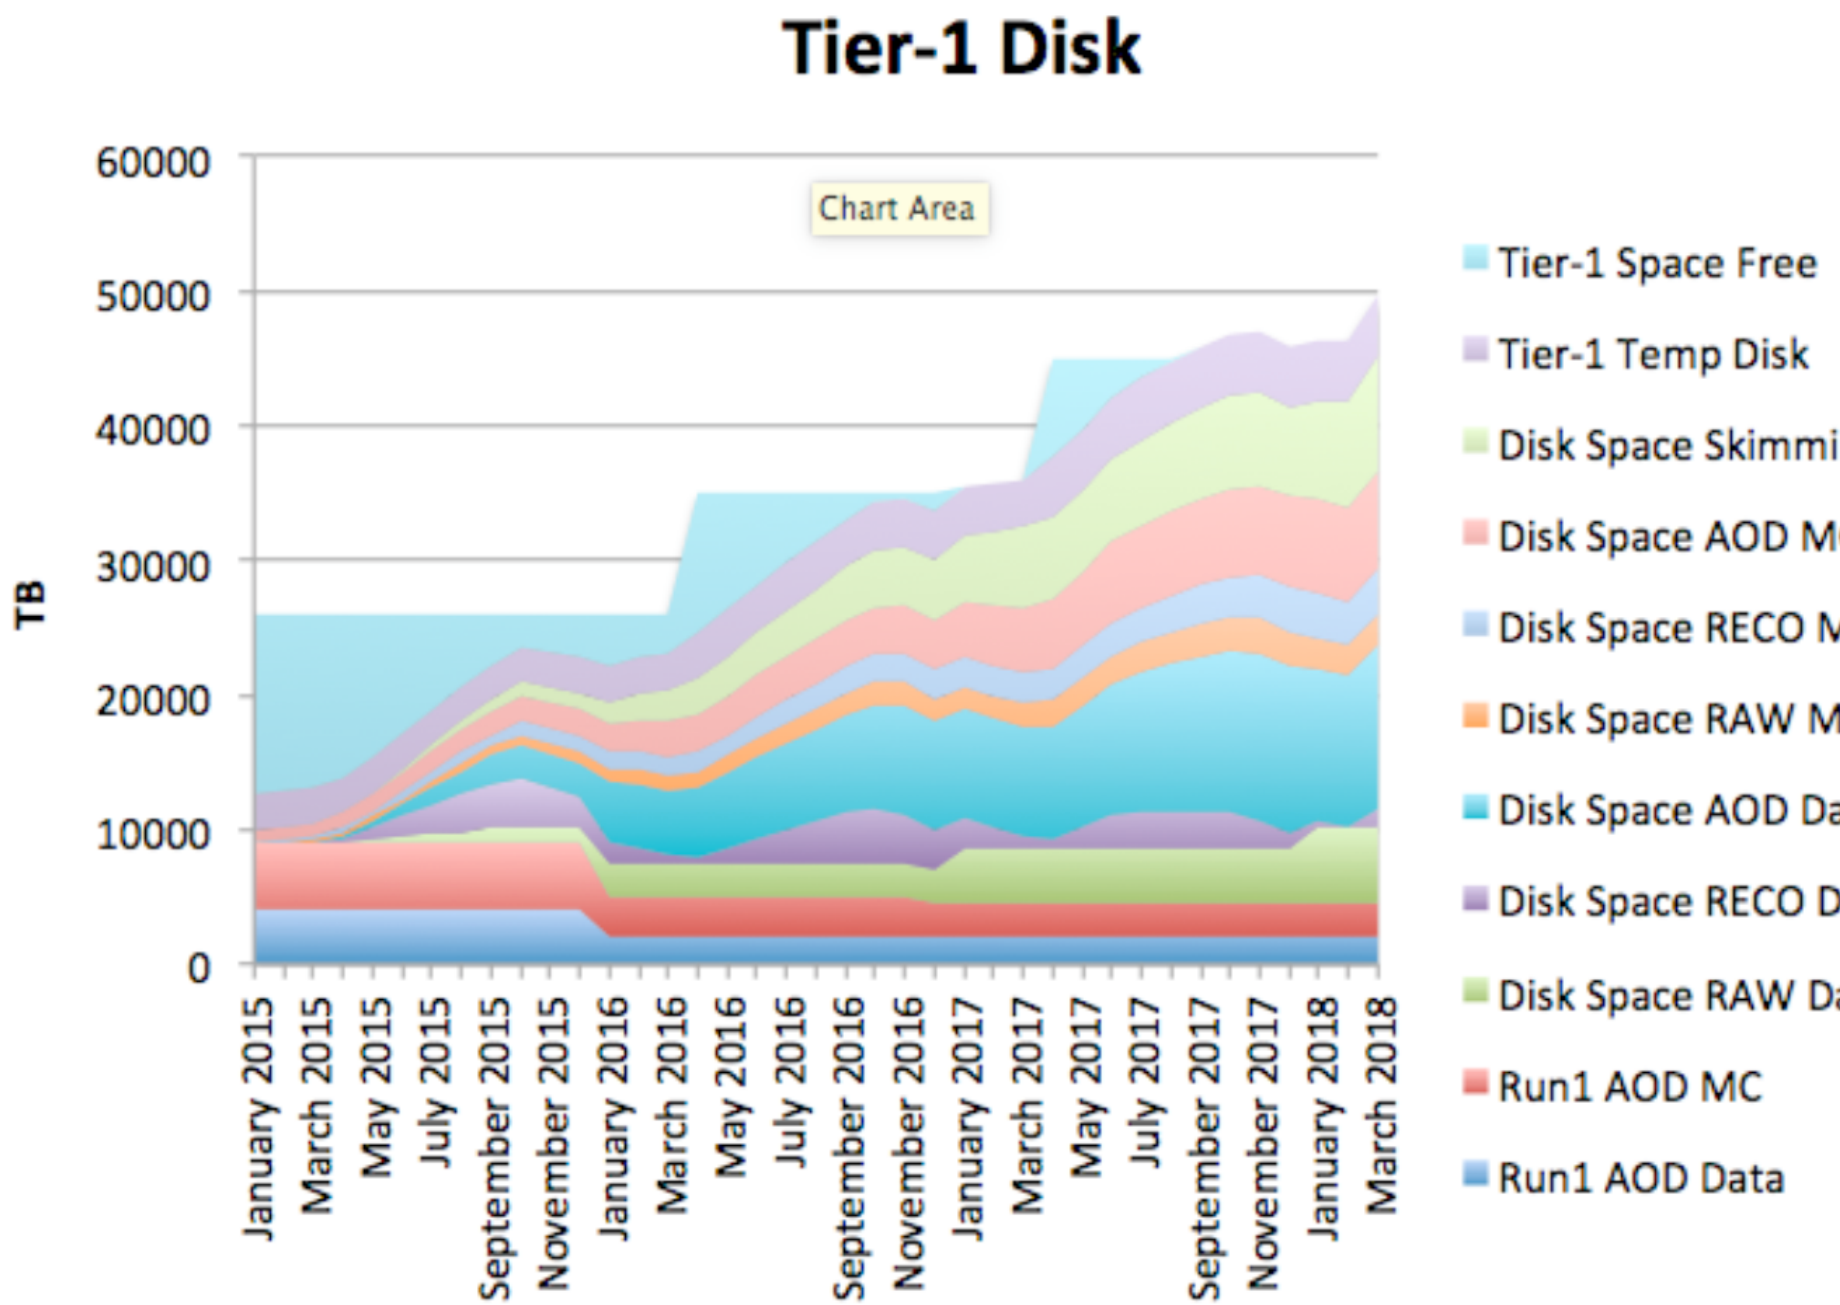

Figure 2: Tier-1 disk usage by data type.

HEP Compute Needs: “spikes” above a large “baseline”

- LHC experiments are looking for ways of fulfilling **peak demands** on the time scale of 2016-2018

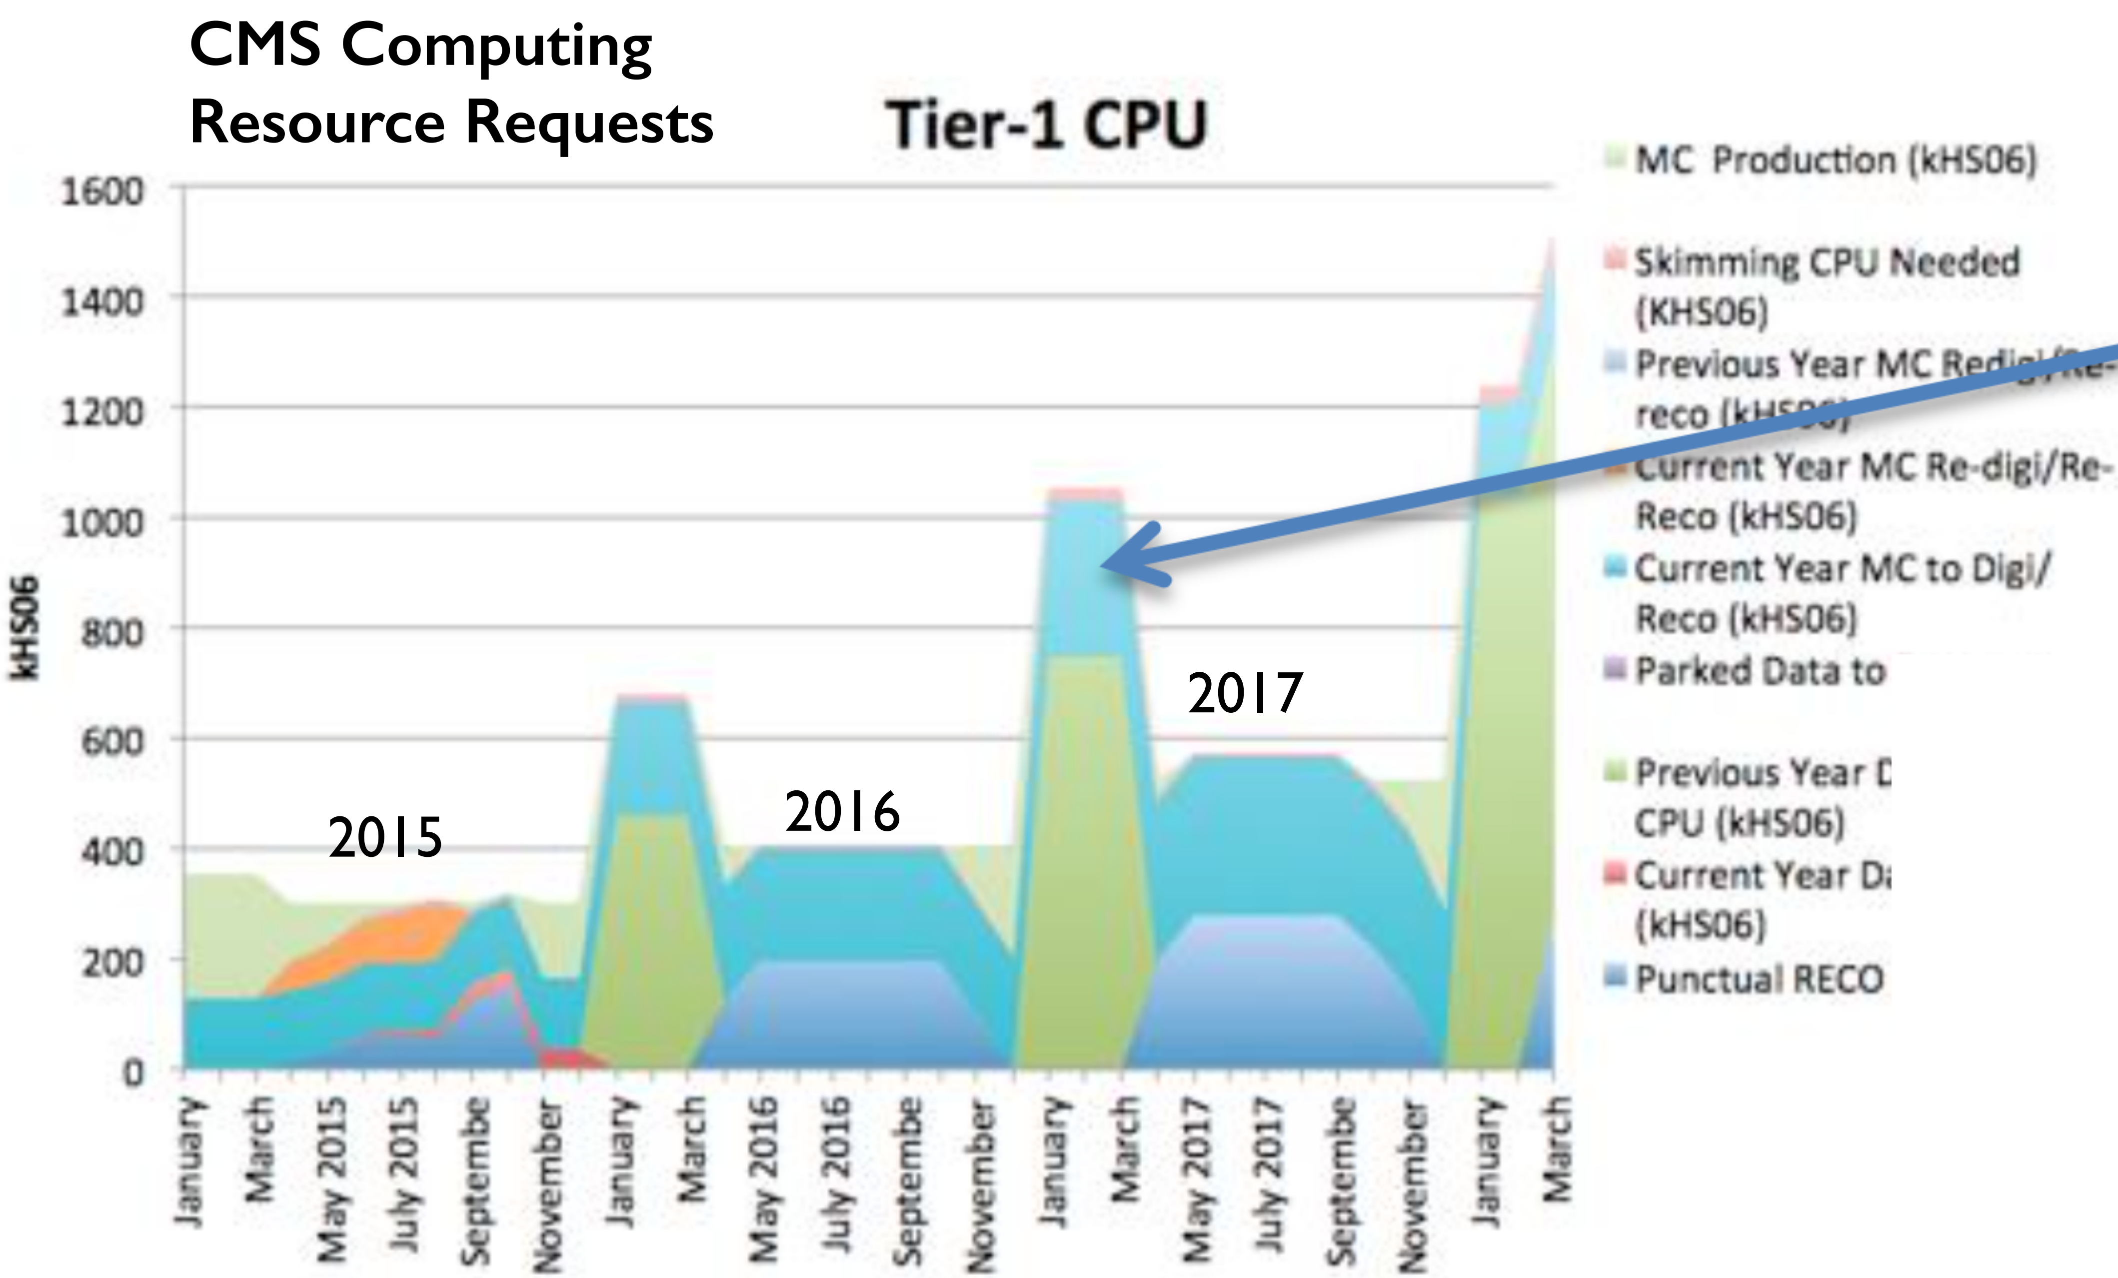

Peaks require on-demand additions to LHC-owned computing infrastructure!

Fermilab

LHC Data Storage Facilities, Future Capacity Needs

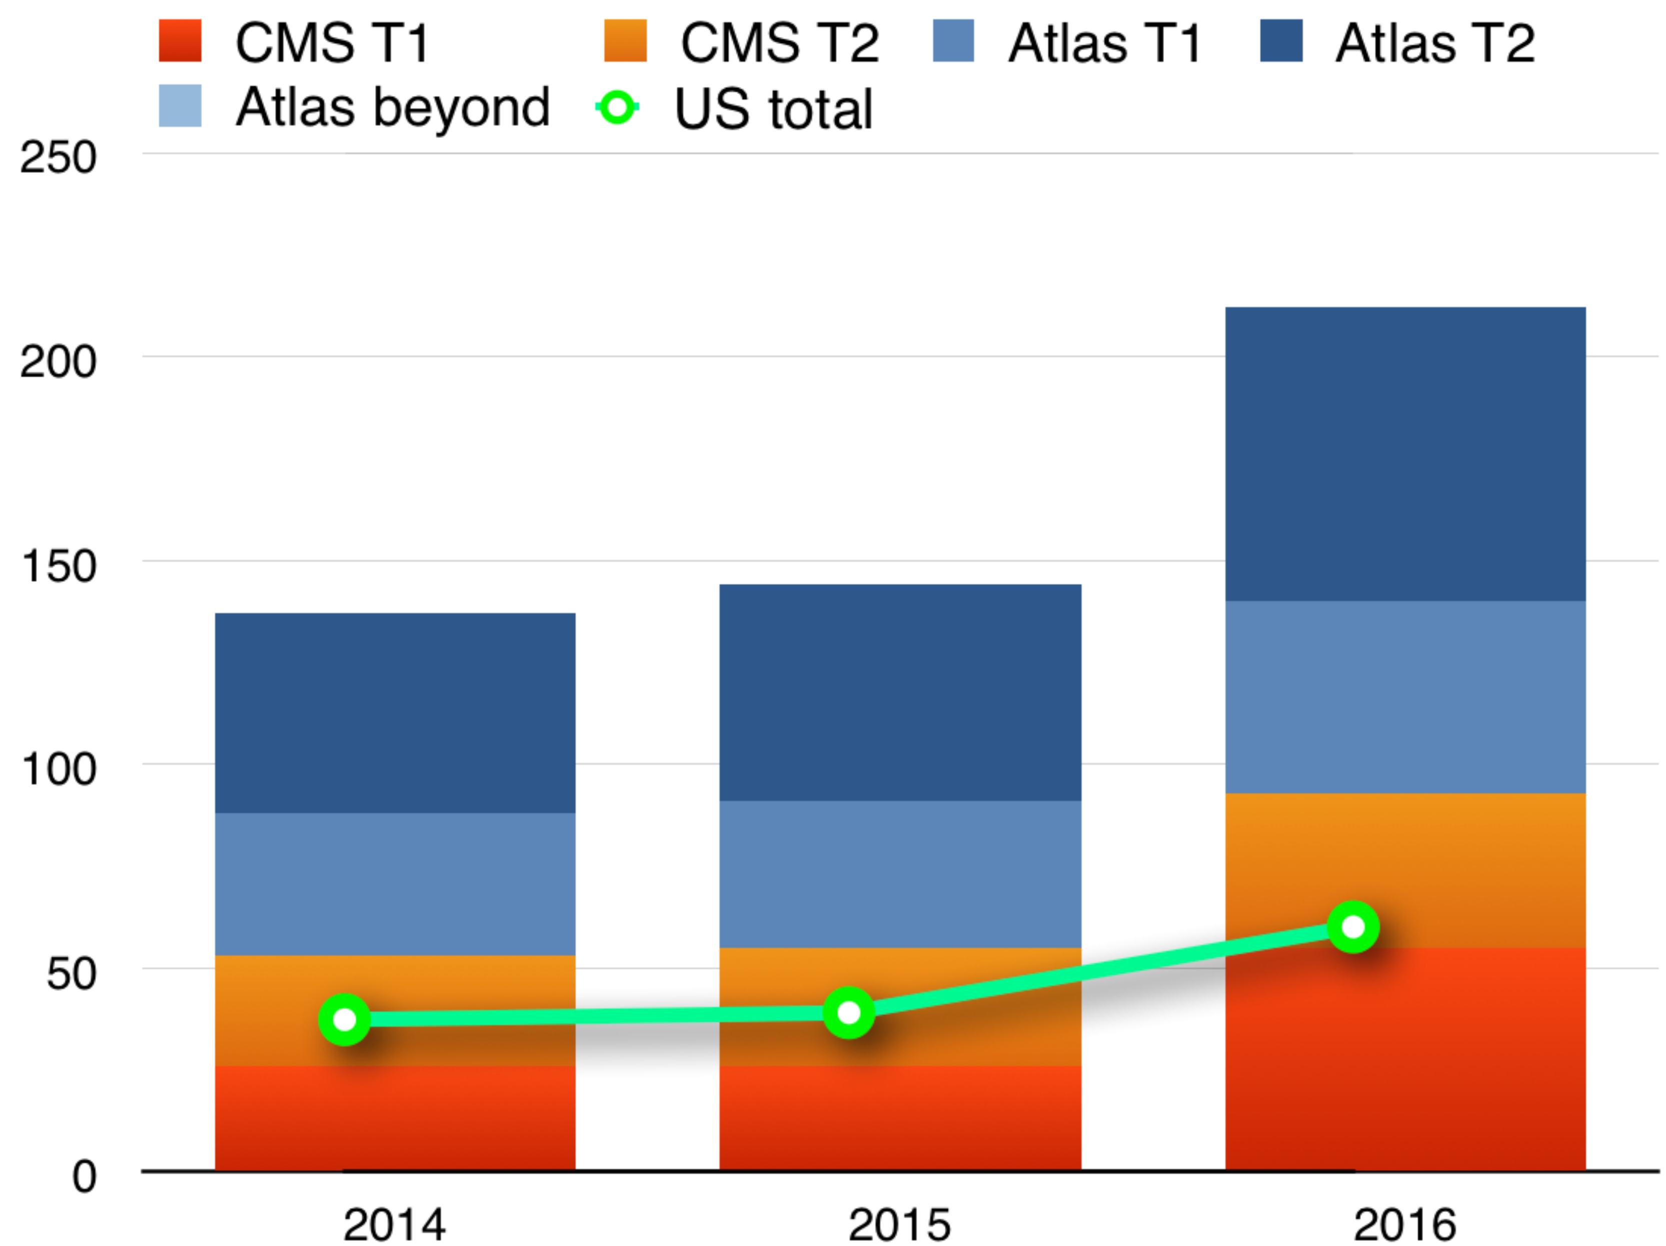

- Disk-based storage systems distributed across the world
- expect at least 60PB in the U.S. in early 2016

LHC Computing Facilities, Future Resource Needs

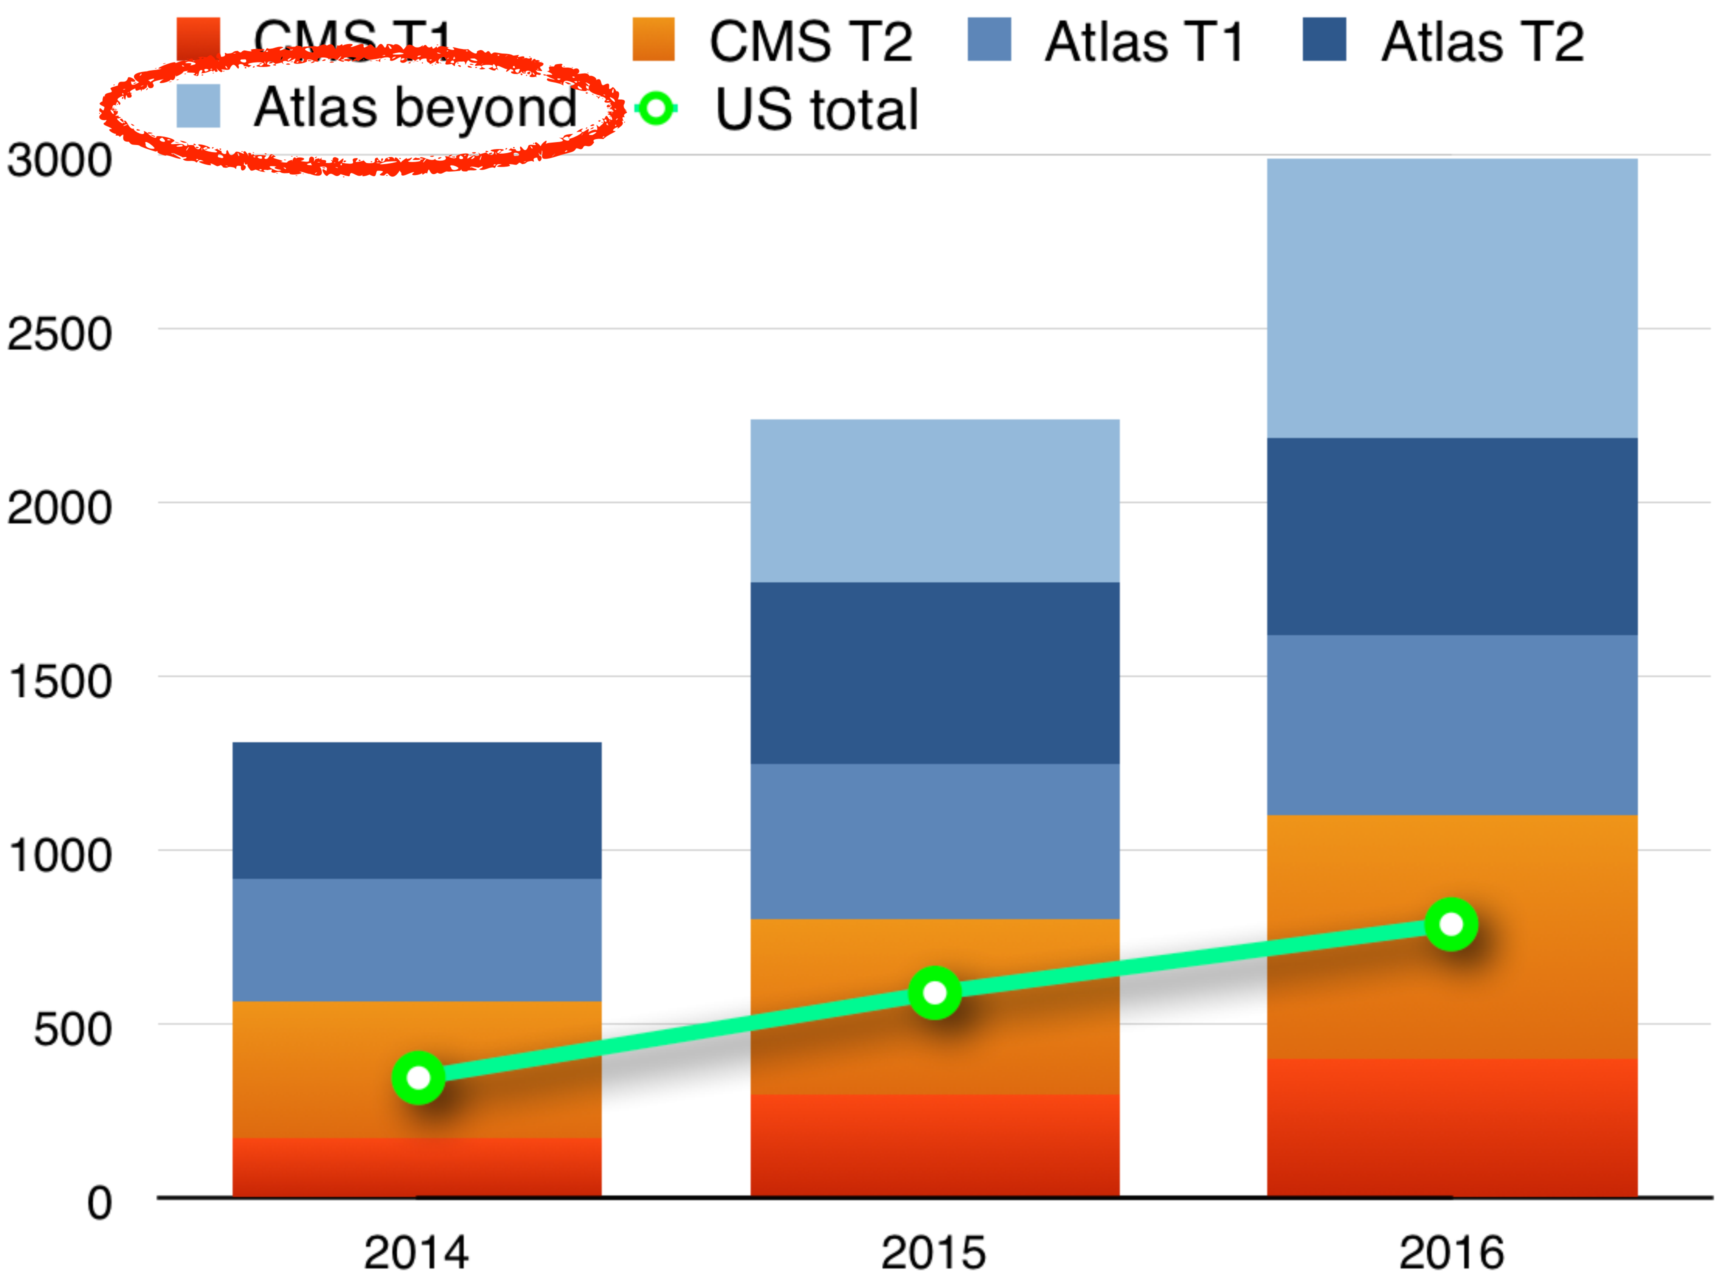

N.B. Atlas model relies on “beyond planned/ pledged” resources ~80k cores in 2016

- Scale: 3,000 kHEPspec06 corresponds to ~ 300,000 cores
- U.S. to provide ~30% of world-wide needs
- These are the “modeled needs” 24x7 — today, CMS has access to ~100,000 cores

Provisioning for Peak Demands

- The "dream" of short turn-around times for workflows
  - ✦ Short latencies in particular in analysis workflows are important for science efficiency
  - ✦ Use resources from a larger pool when they are needed, should also result in more cost-effective solutions
- Separating the processing and storage services allows them to scale independently
- e.g. ATLAS and CMS are looking at ways to double available resources for periods of time
  - ✦ using Amazon services

from M.Girone at CHEP 2015

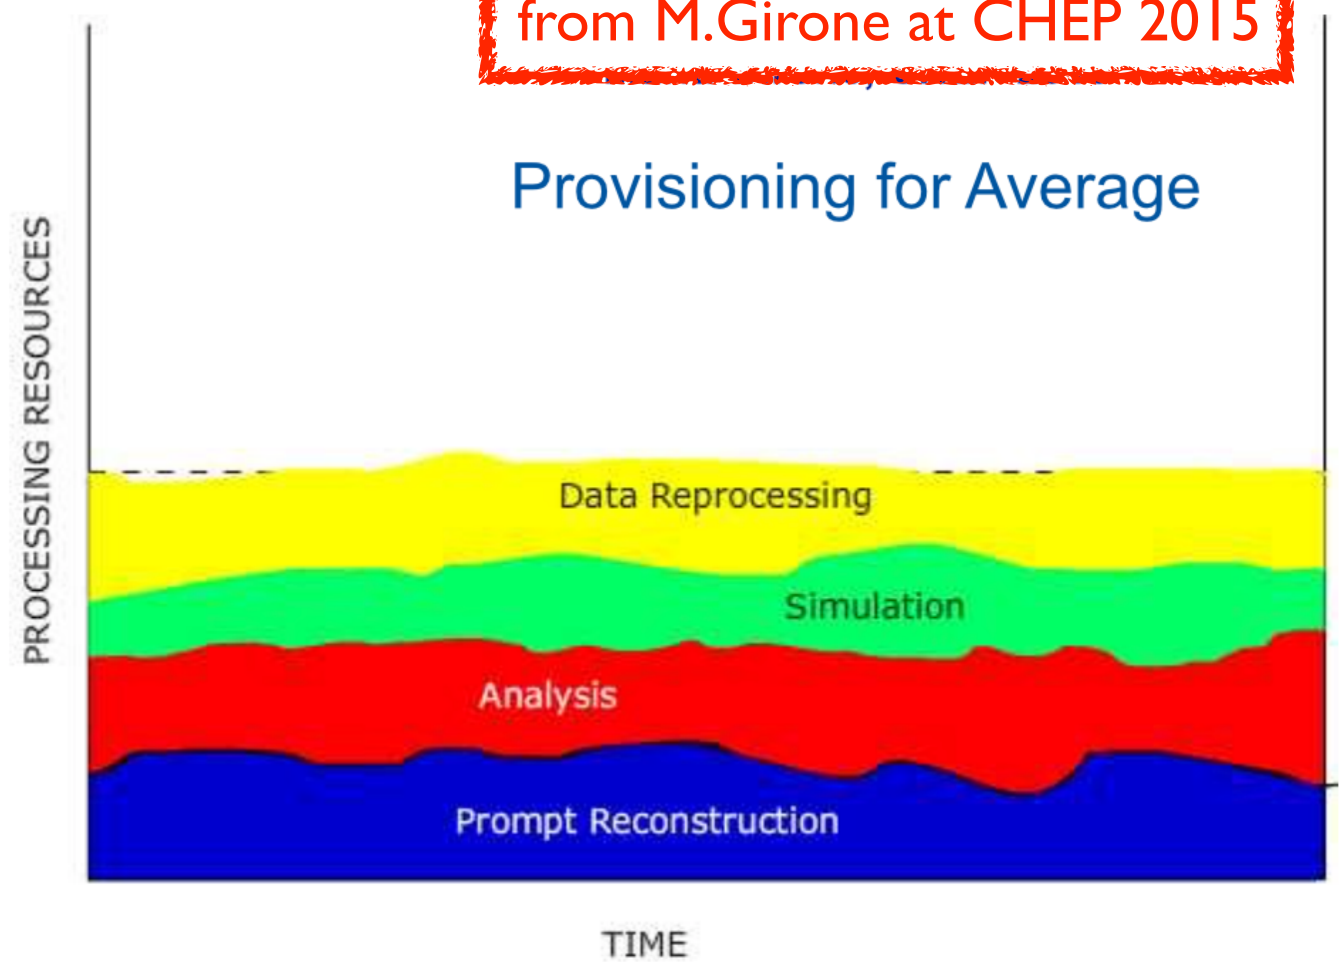

Provisioning for peak requires that we use pooled resource —> Clouds or large HPC Center!

Provisioning for Peak Demands

- The "dream" of short turn-around times for workflows
  - ✦ Short latencies in particular in analysis workflows are important for science efficiency
  - ✦ Use resources from a larger pool when they are needed, should also result in more cost-effective solutions

- Separating the processing and storage services allows them to scale independently

- e.g. ATLAS and CMS are looking at ways to double available resources for periods of time
  - ✦ using Amazon services

from M.Girone at CHEP 2015

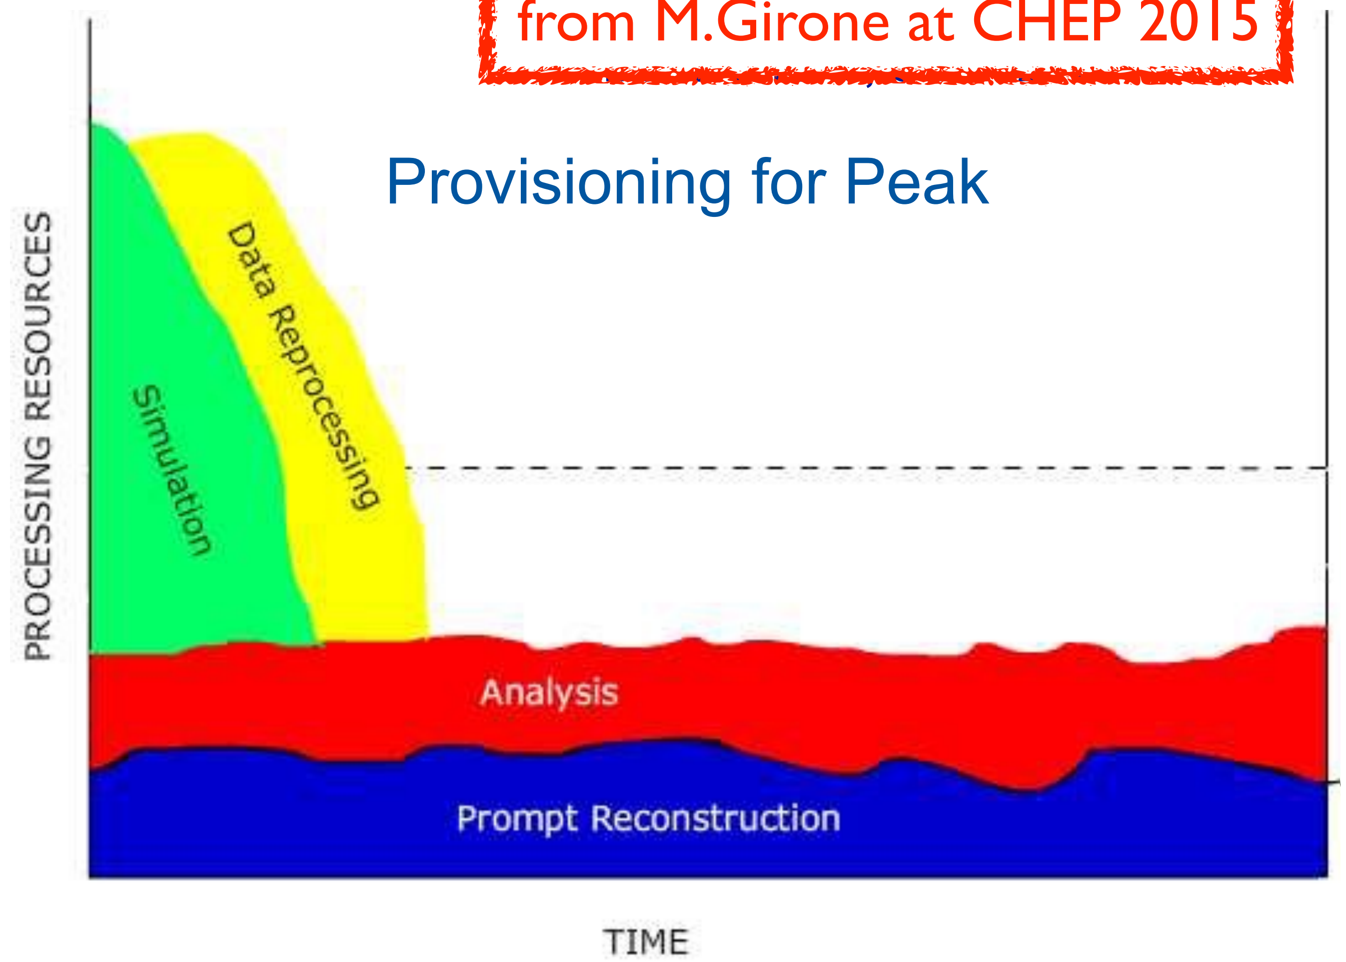

Provisioning for peak requires that we use pooled resource —> Clouds or large HPC Center!

Fermilab

Data access Across the Wide Area Network

- Revolutionizes data management and the provisioning of compute resources!

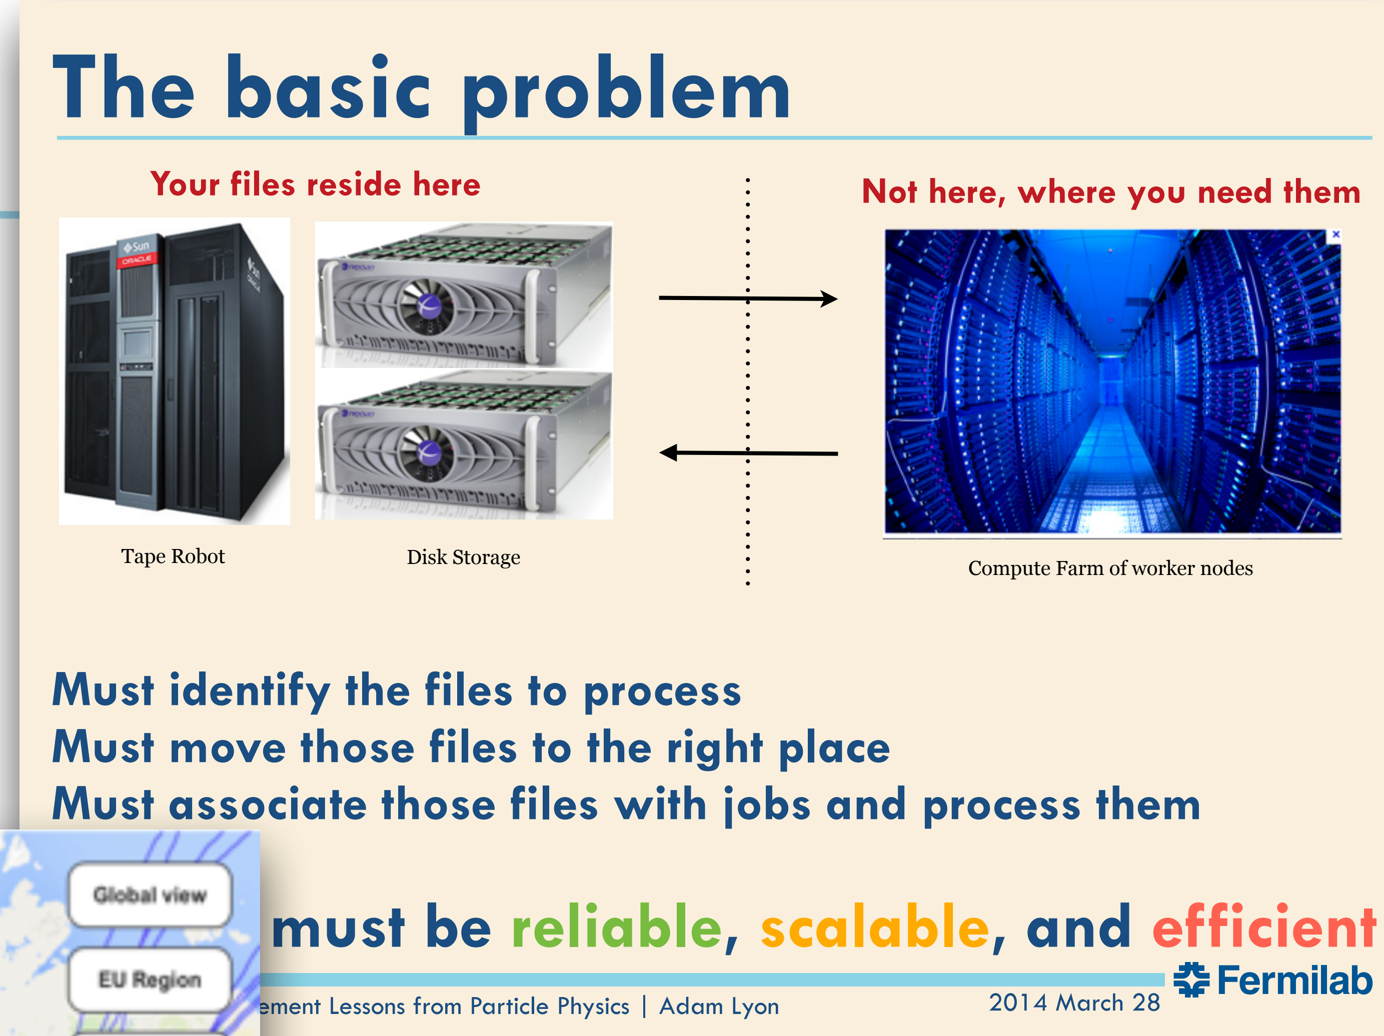

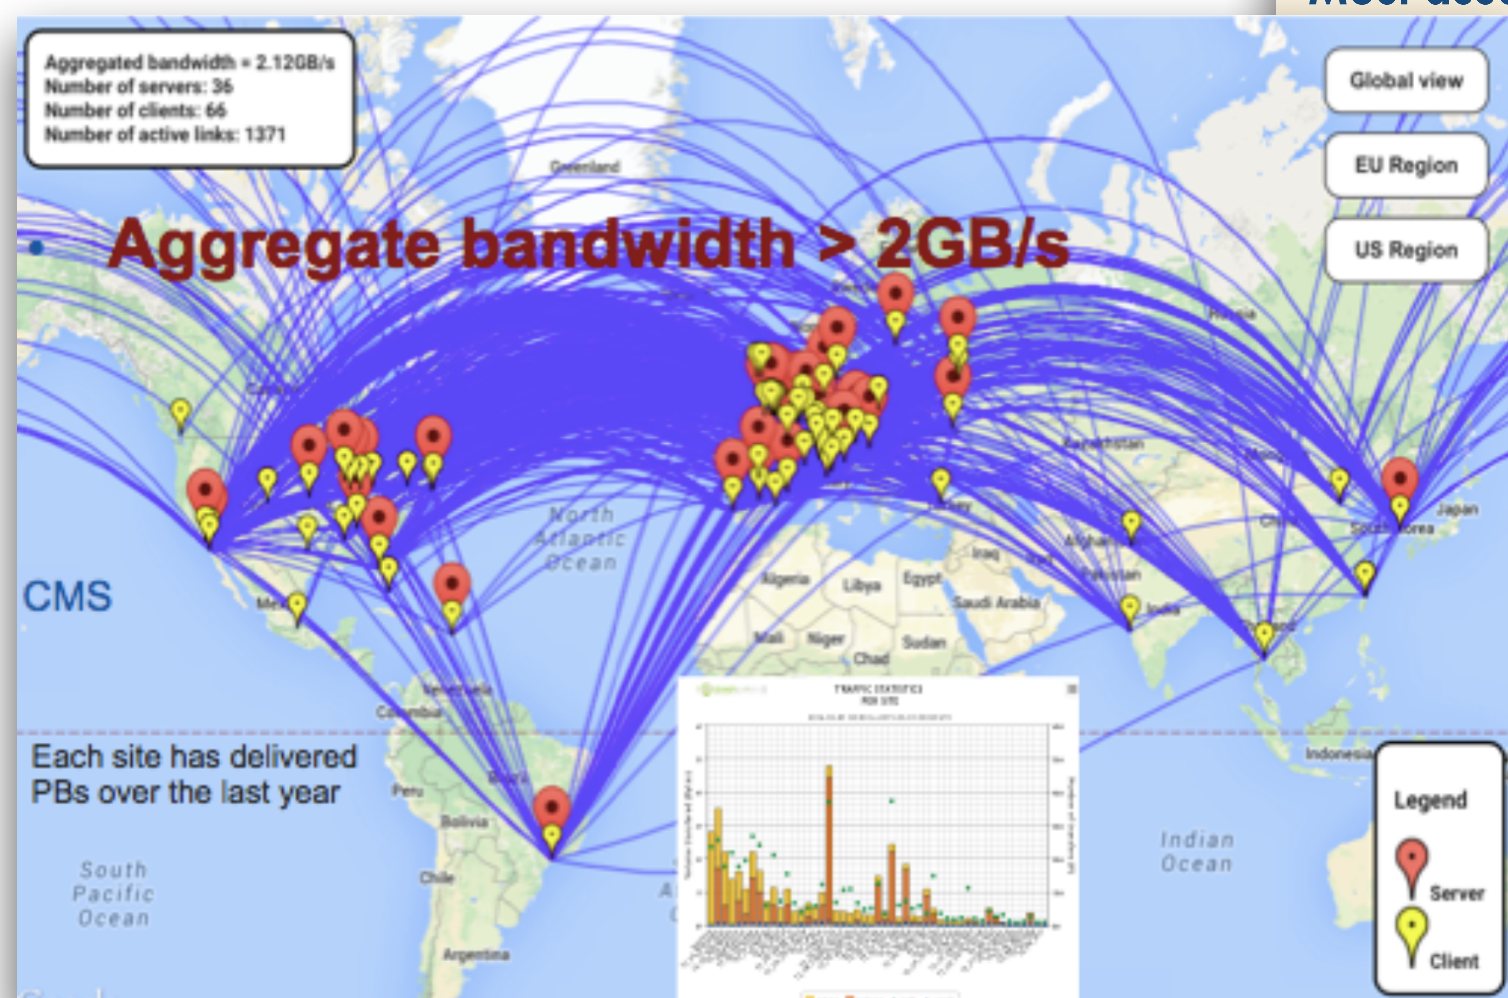

- The WAN *is* reliable, affordable, scalable, and efficient!
- Security aspects have been addressed
  - traceability and evolving trust relationships model

This has huge impact on how science communities can use any kind of compute resources today!

O.Gutsche's talk at CHEP 2015

This talk

- In the recent past, HEP resources were firmly based on **Grid** technologies
  - HEP applications == HTC
    - High Throughput Computing applications

- The need for more **capacity** and **elasticity** makes us look at other resource providers:
  - **Cloud**
  - **HPC** = High Performance Computing

GRID

Cloud

HPC

Fermilab

4 Gabriele Garzoglio, Oliver Gutsche I CHEP2015: Diversity in Computing Technologies and Strategies for Dynamic Resource Allocation 15. April 2015

This opens up the “phase space” of available and affordable resources

10,000 feet overview

O.Gutsche’s talk at CHEP 2015

Grid

- Virtual Organizations (VOs) of users trusted by Grid sites
- VOs get allocations ➜ **Pledges**
  - Unused allocations: opportunistic resources

Trust Federation

Cloud

- Community Clouds - Similar trust federation to Grids
- Commercial Clouds - **Pay-As-You-Go** model
  - Strongly accounted
  - Near-infinite capacity ➜ **Elasticity**
  - Spot price market

Economic Model

HPC

- Researchers granted access to HPC installations
- Peer review committees award **Allocations**
  - Awards model designed for individual PIs rather than large collaborations

Grant Allocation

Fermilab

7 Gabriele Garzoglio, Oliver Gutsche I CHEP2015: Diversity in Computing Technologies and Strategies for Dynamic Resource Allocation 15. April 2015

For example, use of commercial clouds becomes viable

M.Girone's talk at CHEP 2015

Integrating Clouds

4

- Amazon is huge potential resource
  - They have **56M** processor cores, which dwarfs our 100k
  - Commercial cloud services are typically very expensive, but with these grants we cut the cost by a factor of ten

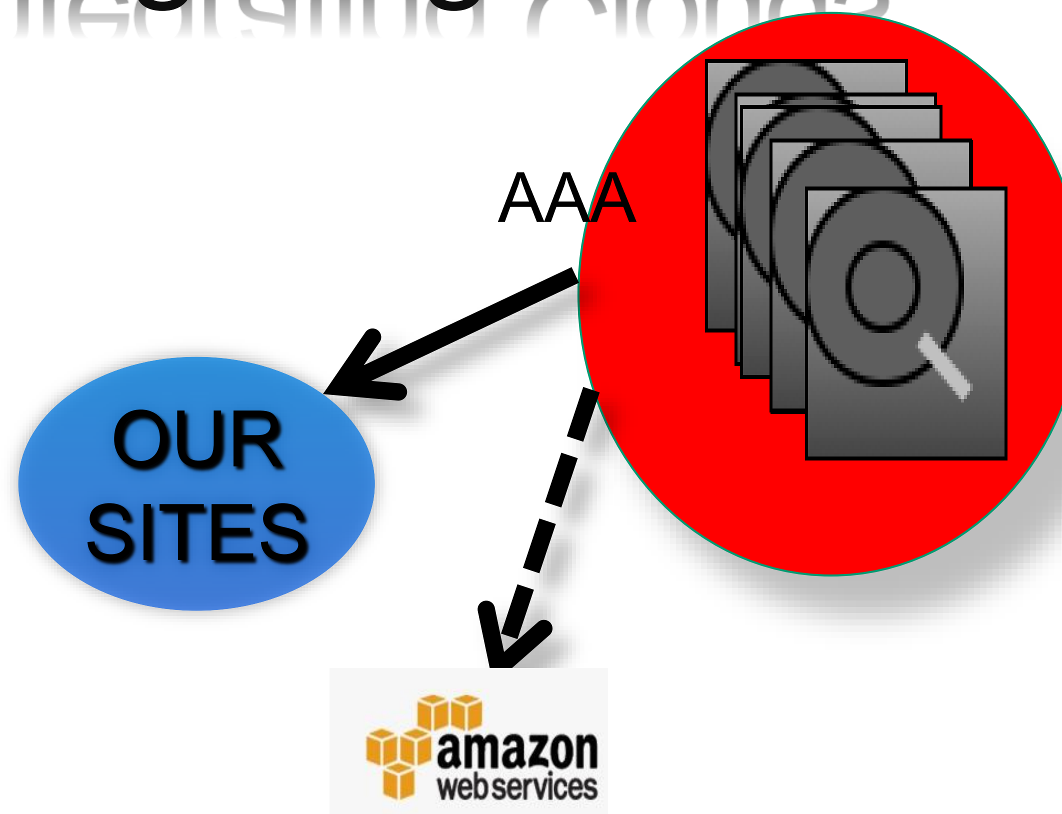

- The cloud grants allow us to test a new method of provisioning resources
  - Processing is very large scale, local storage is transient, and the bulk of the data is served over the federation
  - Opens a new class of computing

Maria Girone, CERN

This new paradigm significantly changes the role of Facilities in providing end-to-end solutions to their customers

- Facilities can benefit from providing a much more “elastic” offering
- Moving away from stand-alone stove-piped facilities toward Labs providing leadership role in the computing **eco-system**
  - labs, leadership class and production class facilities, etc are part of and leaders in the eco system, that includes scientific and commercial providers
  - sites will find a large “market” for specific offerings: specialized architectures, archival capabilities, database services, data management solutions
  - the eco-system is enabled by the labs and OSG and others
- Facility’s role is still to provide “complete solutions” for their users
  - CPU and data capacities with guaranteed level of service
  - Users would not have to care about wether their jobs are running on “owned” or “rented” resources Sites could make the economic decision themselves and optimize their cost structure
  - Storage services that adapt to where the jobs are running
  - On-demand services that scale by tapping into large pooled resources
    - like clouds, HPC, OSG etc

Fermilab is moving ahead with the HEPCloud Facility To provision local, cloud and HPC resources to EF and IF

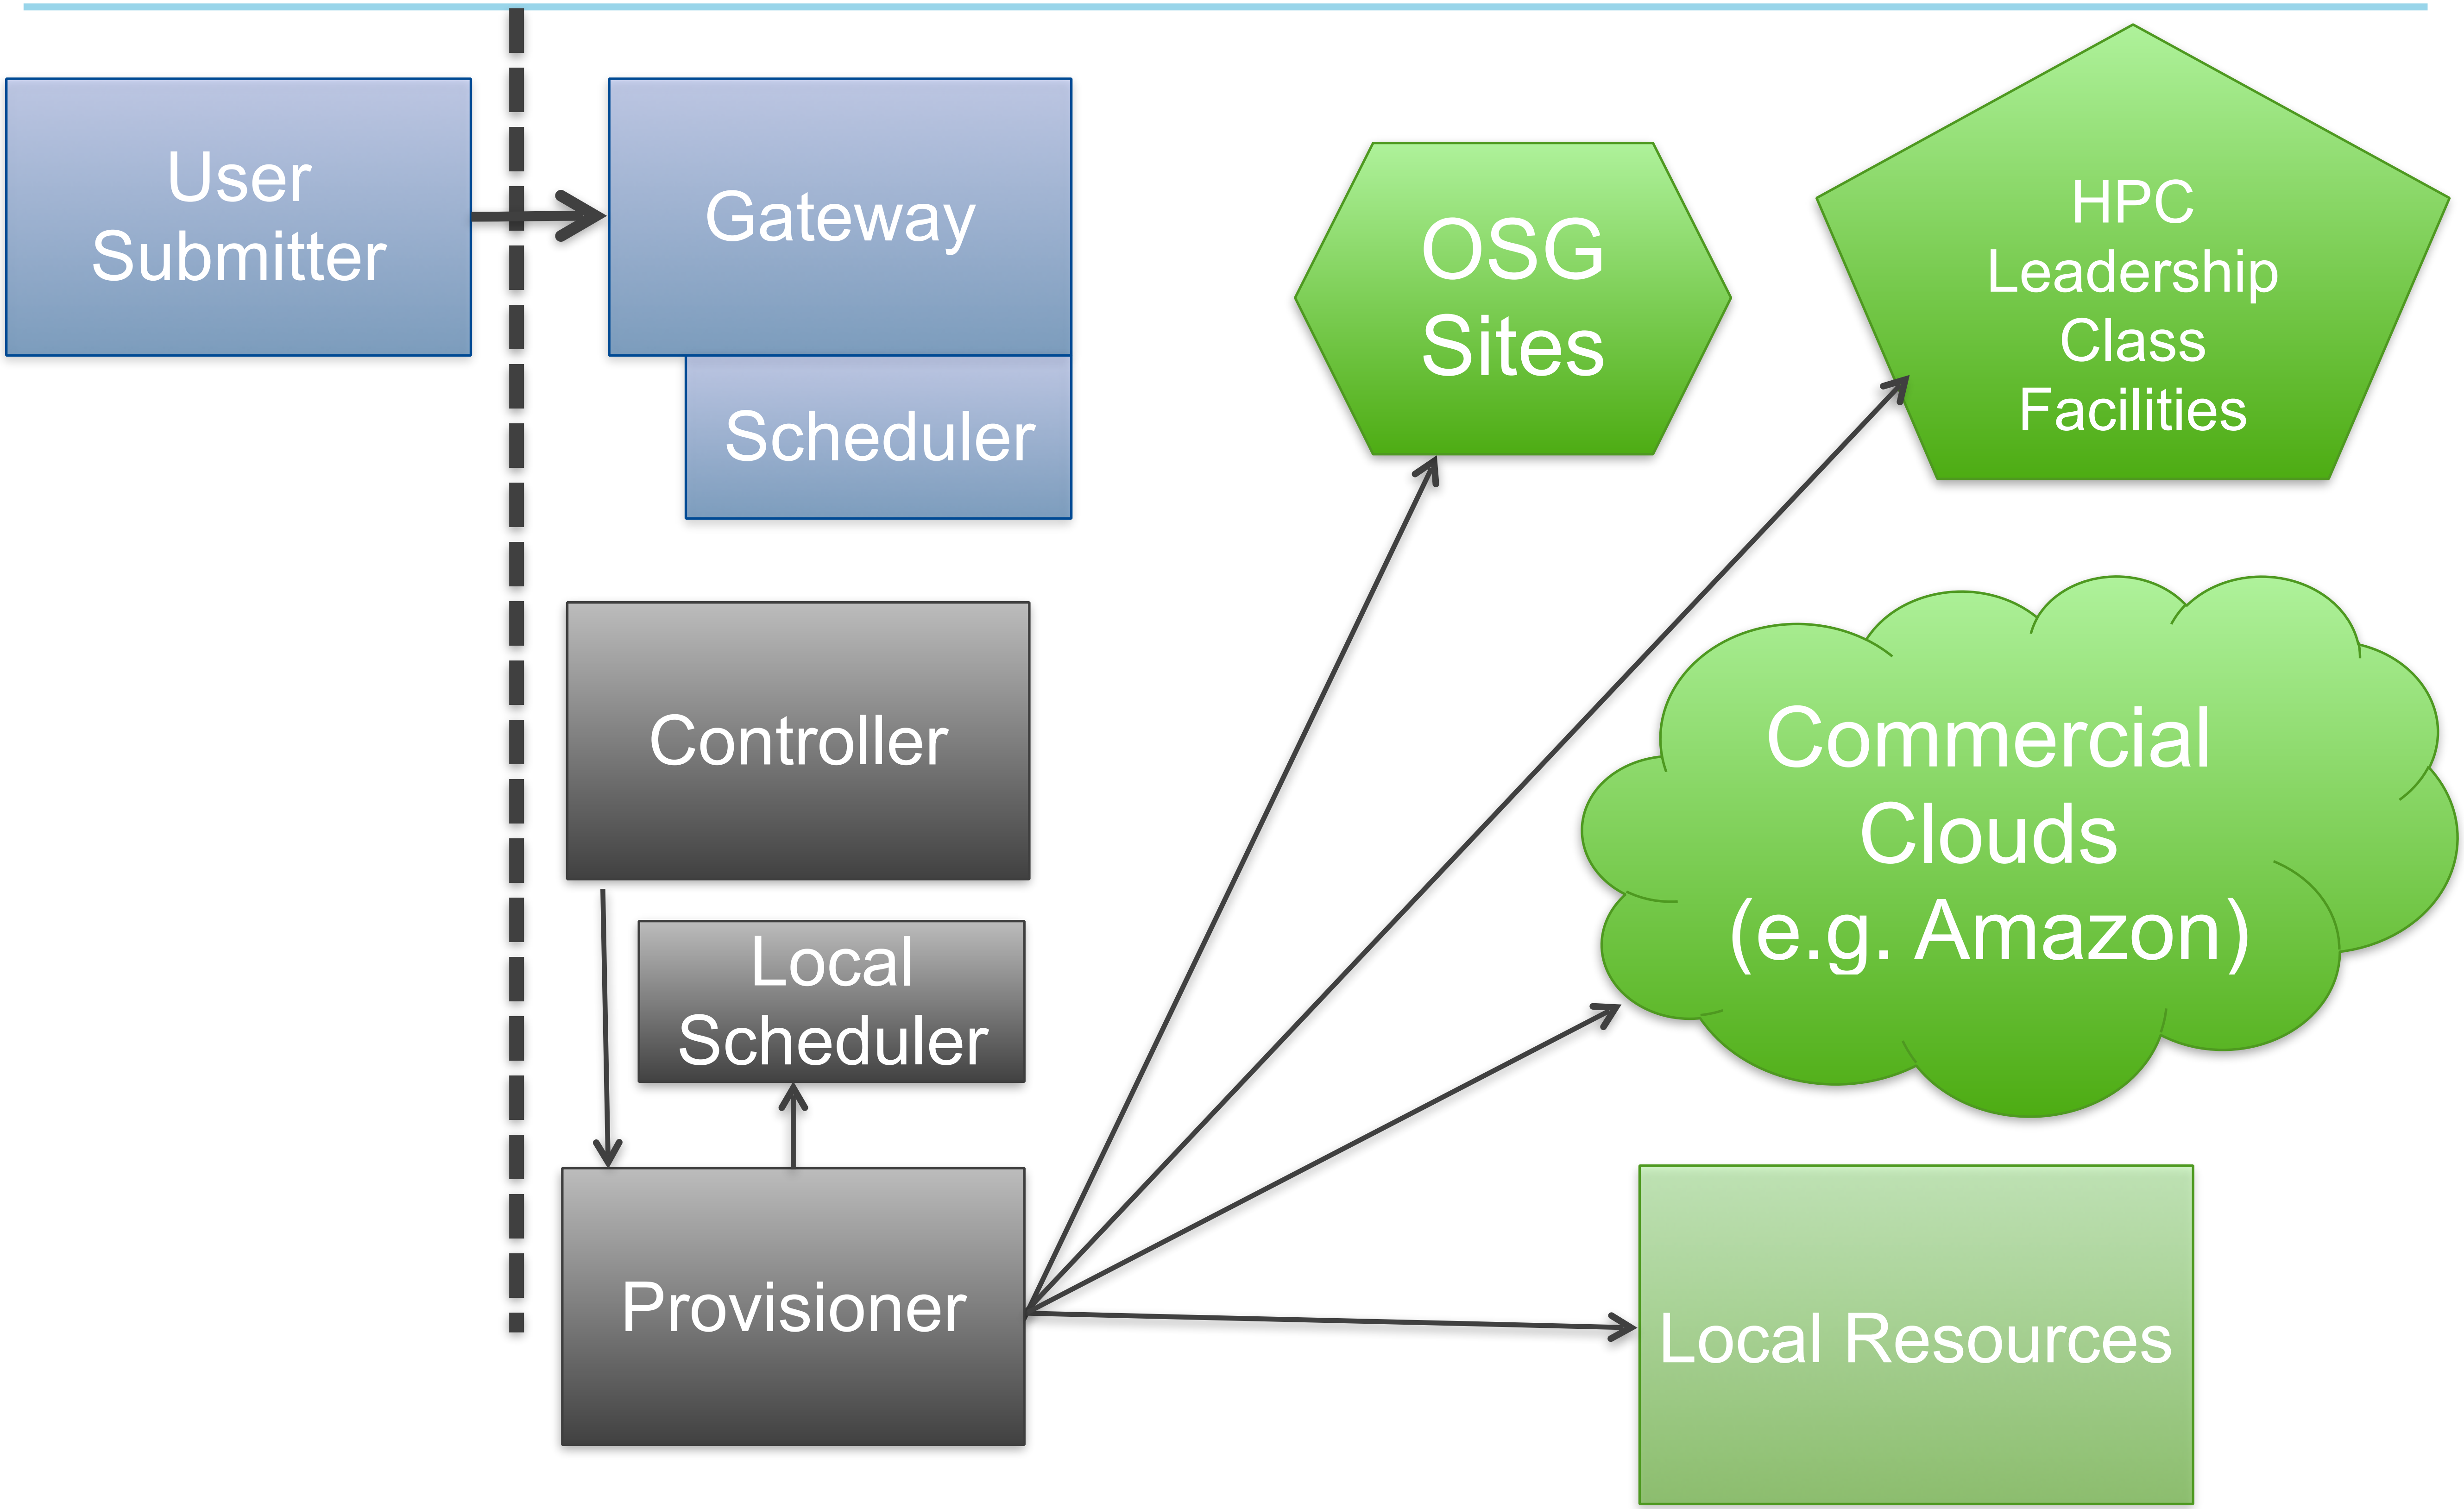

ESnet provides Trans-Atlantic Connectivity for HEP as a World-Wide Effort

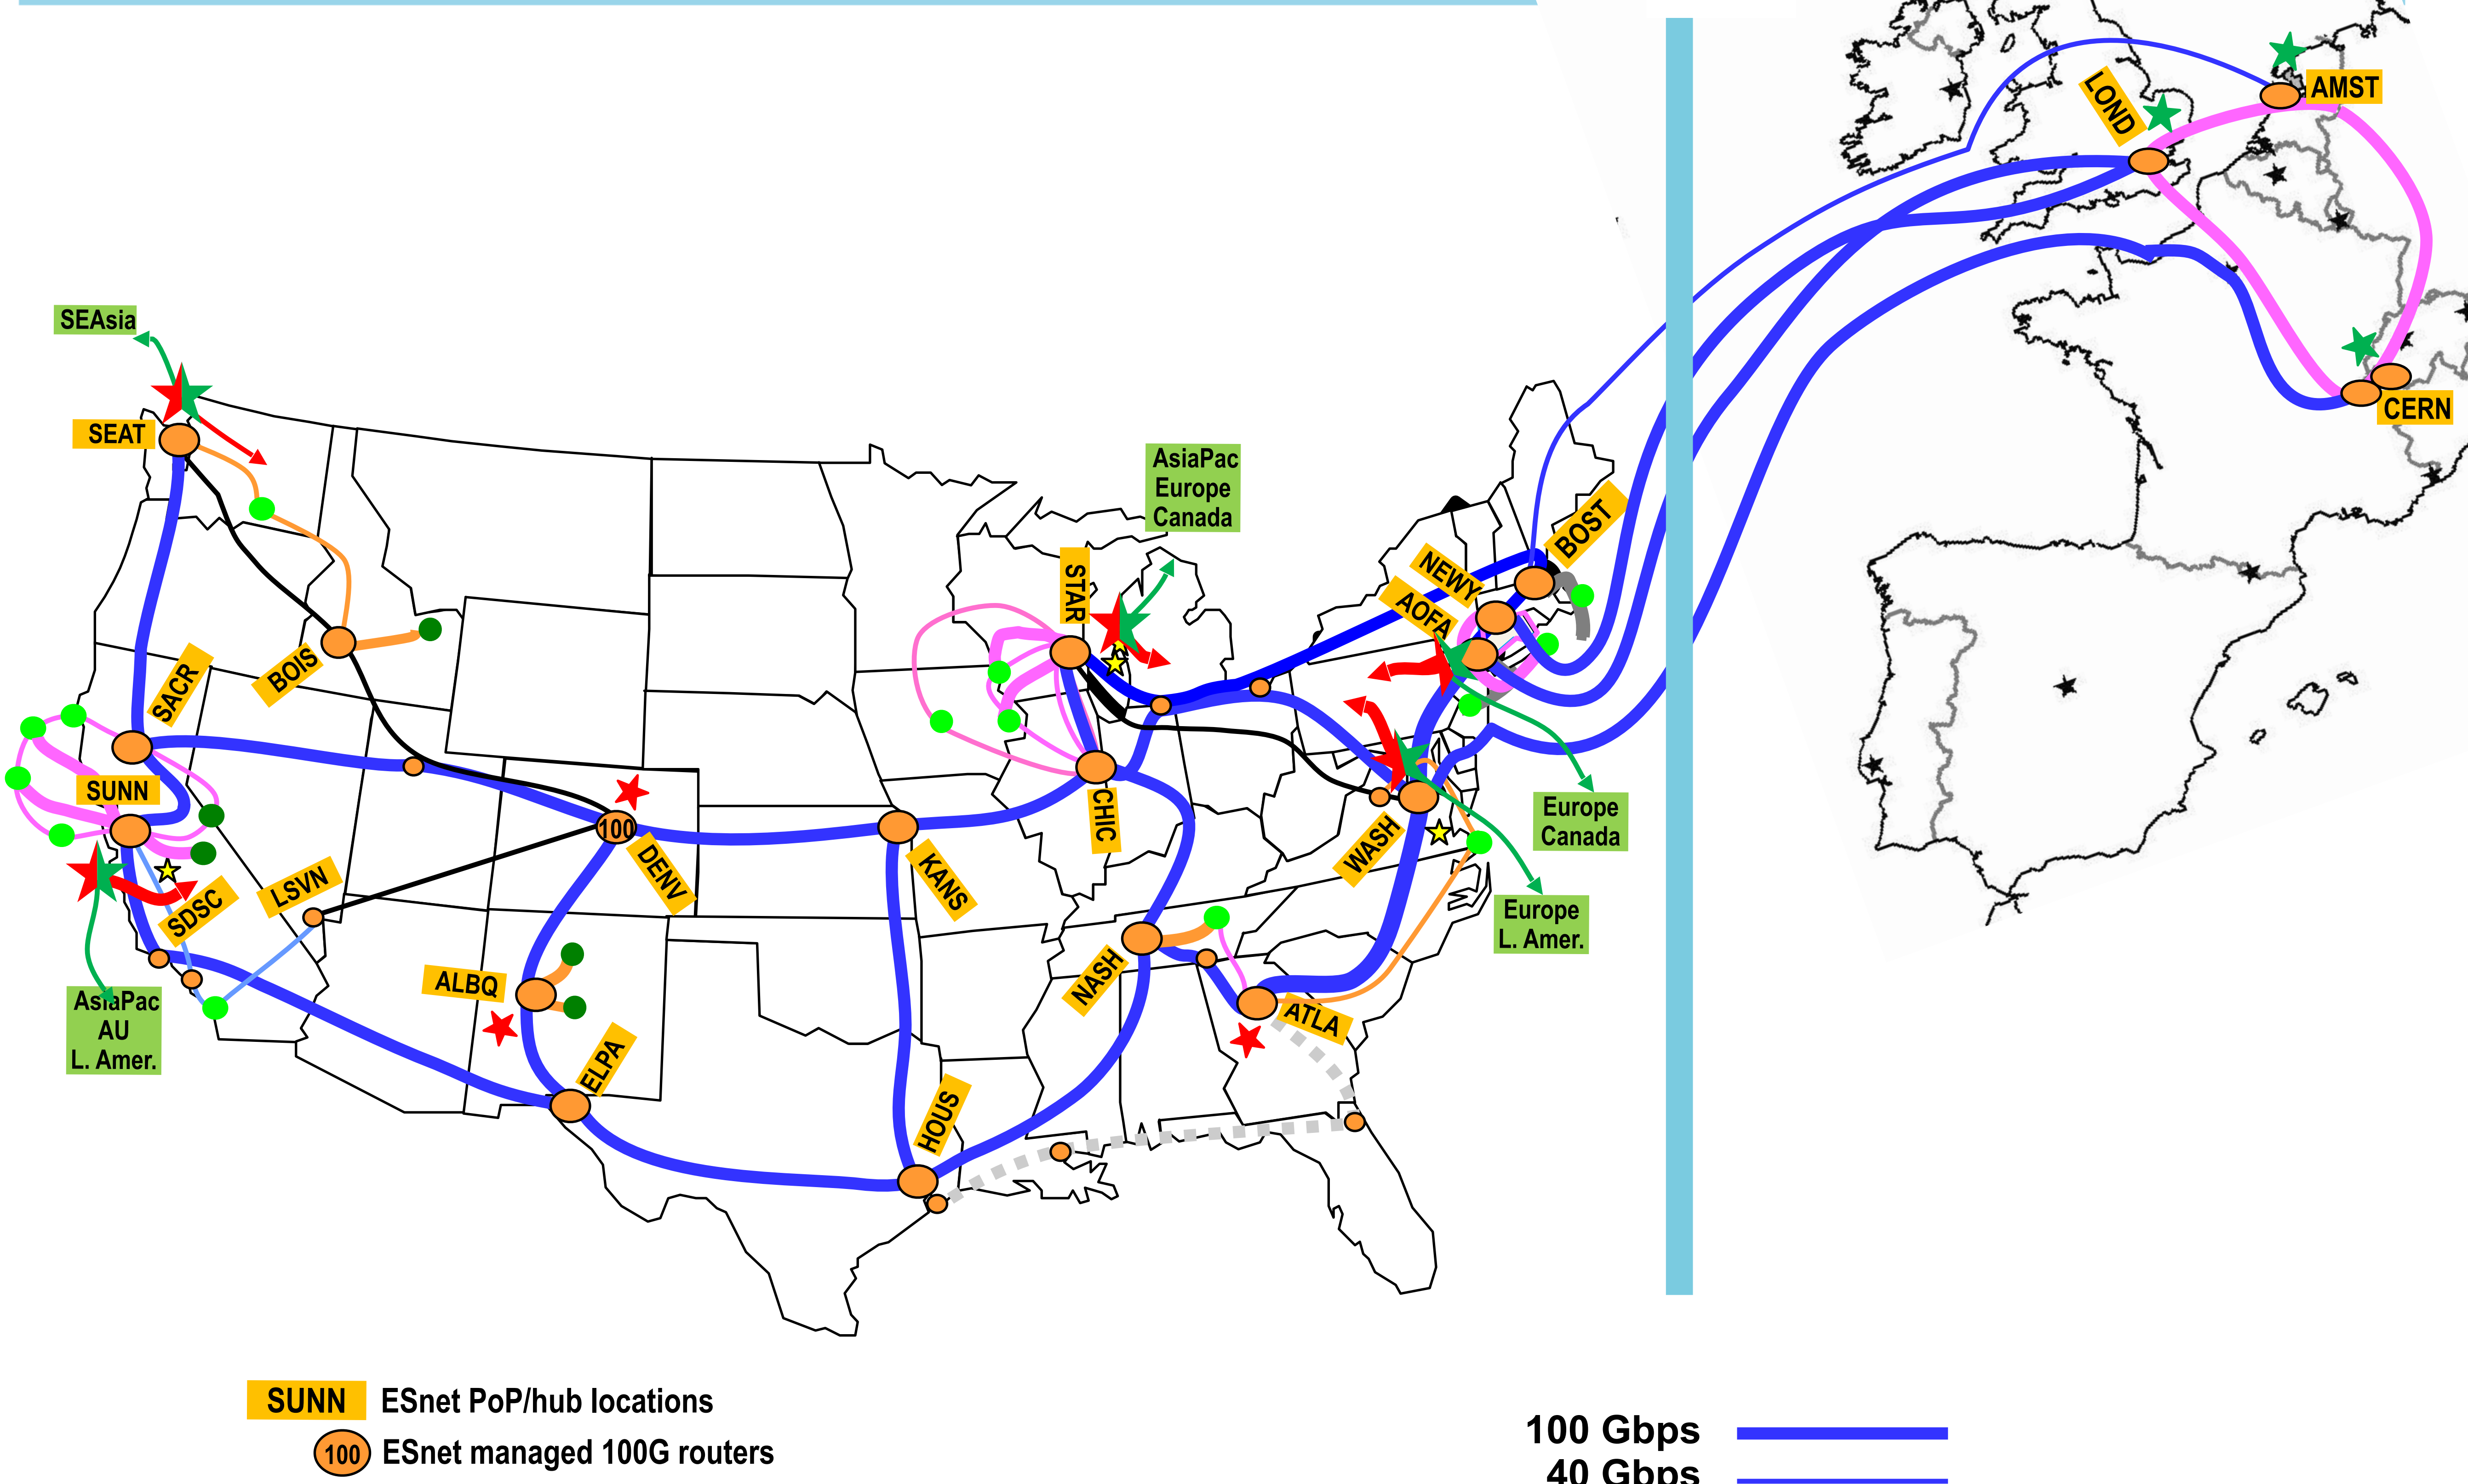

Fermilab

A new Era of Cooperation and Collaboration in Computing, Between Labs and Universities, Between Sciences, International

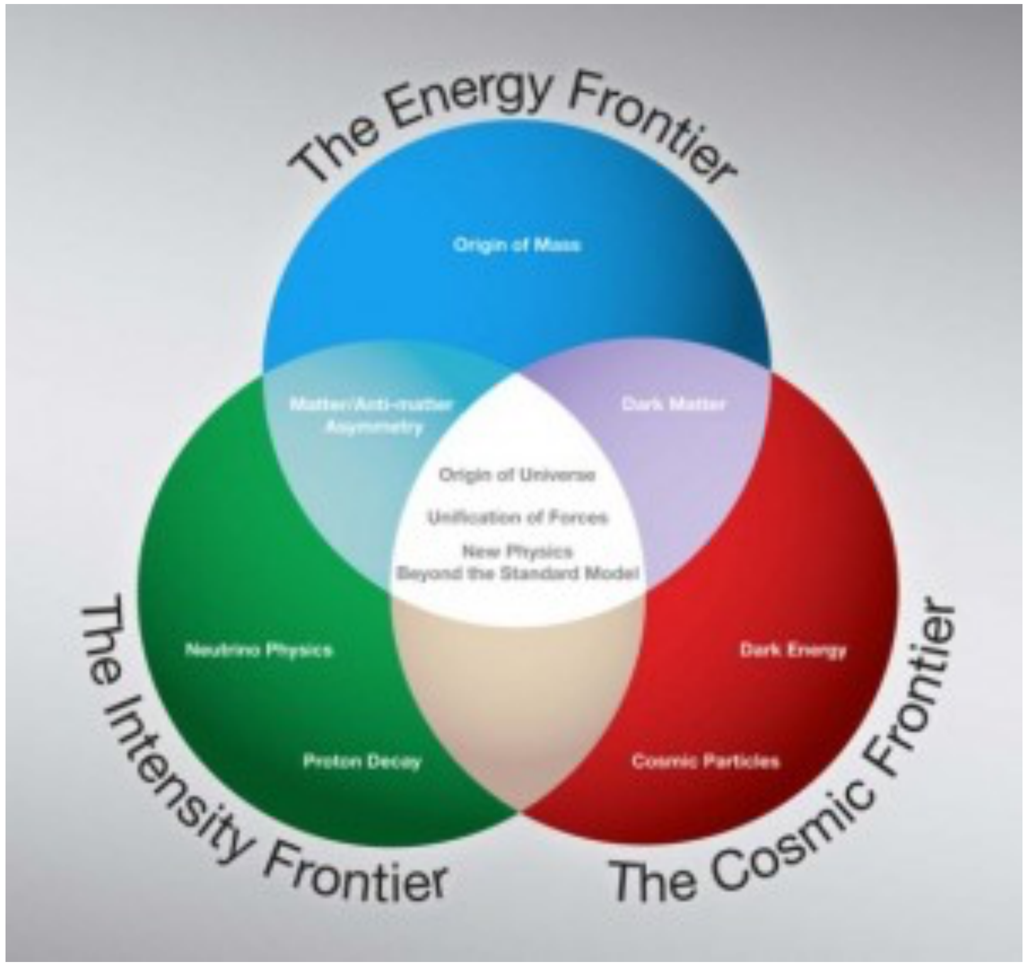

- Very concrete examples for cooperations in all directions
  - like the ESnet Extension to Europe EEX supporting LHC TA network
  - the HEP-FCE Forum for Computational Excellence
  - the HEP Software Foundation with Europe, many more

Summary and Outlook

- There is no end in sight for large increases in resource demands and new capabilities, which change expectations and requirements on HEP facilities
  - provide services to distributed communities, supporting complex end-to-end use cases involving huge computational and data throughput needs and capabilities
- The role of the facility providers are changing as they are facing cost effective competition to their “bare metal” offerings from IaaS providers
  - Facilities remain to be first-line support for the complex scientific work flows and data management needs of HEP and other DOE SC communities
  - Facilities should integrate new opportunities and capabilities into their service offerings, in particular in connecting to large data management and data access systems, beyond “login and batch” services for applications and application libraries
- Facilities should keep an open mind how to provide their services so they fit into and enrich the US and international scientific computing eco-system
  - requires new thinking and approaches to difficult issues in the distributed environment, including security, robustness and protection of resources, accounting, prioritization etc
  - good experiences with LHC and emerging IF experiments
- HEP and other DOE Facilities clearly have a huge opportunity for great leadership roles in this environment

Backup Slides

LHC Upgrades and High-Luminosity LHC

Long Shutdowns

LS1 2013-14

LS1 : Prep CMS for ~13 TeV, >1 x 10^{34} Hz/cm^2, <PU> >25, 25ns bunch spacing

- 4^{th} Endcap Muon station, improve readout of CSC ME1/1 & DTs
- Replace HCAL HF and HO photo-detectors
- Tracker operation at -20◦C
- Prepare for further Phase 1 upgrades

Run 2

LS2 2019

Phase 1 upgrades: Prepare for 1.6 x 10^{34} Hz/cm^2, <PU> ~40, ≤200 fb^{-1} by LS2,
and up to 2.5 x 10^{34} Hz/cm^2, <PU> ~ 60, <500 fb^{-1} by LS3

- New L1-trigger system (Calorimeter - Muons - Global) (ready for physics 2016)
- New Pixel detector (installation in technical stop, start of 2017)
- HCAL upgrade: photodetectors and electronics

Run 3

LS3 2023-25

Phase 2 upgrades: ≳ 5 x 10^{34} Hz/cm^2 luminosity leveled, <PU> ≳ 140
Reach total of 3000 fb^{-1} in ~10 yrs operation

- Replace detector systems whose performance is significantly degrading due to radiation damage
  - Tracker (pixels and strips), Endcap calorimeters
- Maintain physics performance at this very high PU
  - Trigger, electronics, enhanced tracker coverage

- 3000 fb^{-1} at $\sqrt{s}$=14 TeV
  - **Precision studies of Higgs and any other new particles yet to be found**

New physics reach example

- SUSY EWK gauginos

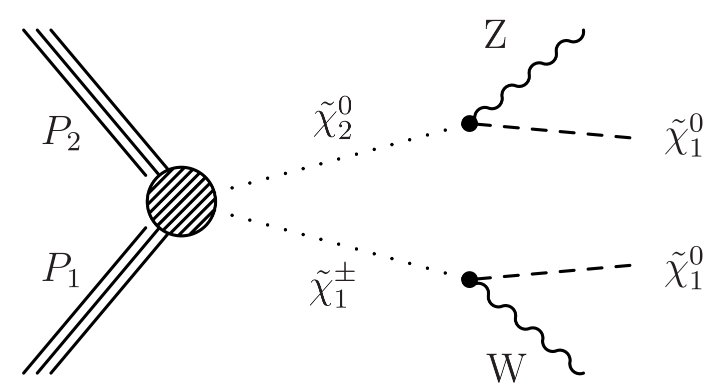

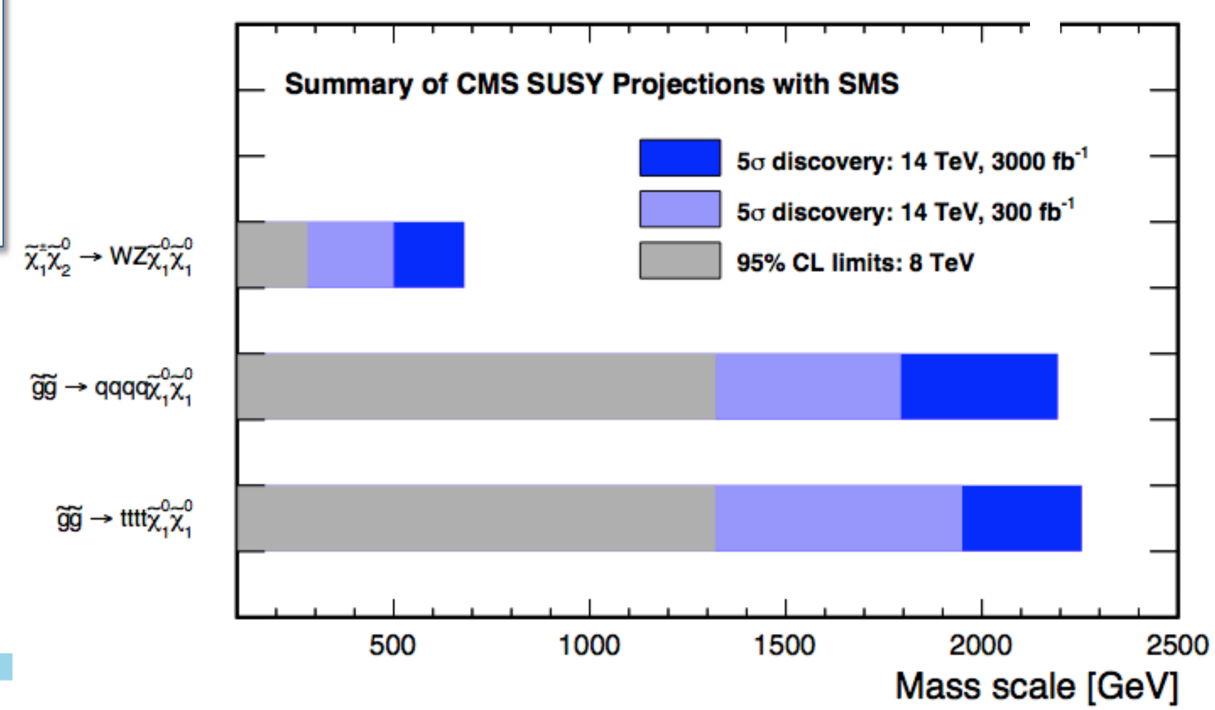

from G.Steward

LHC Machine Evolution

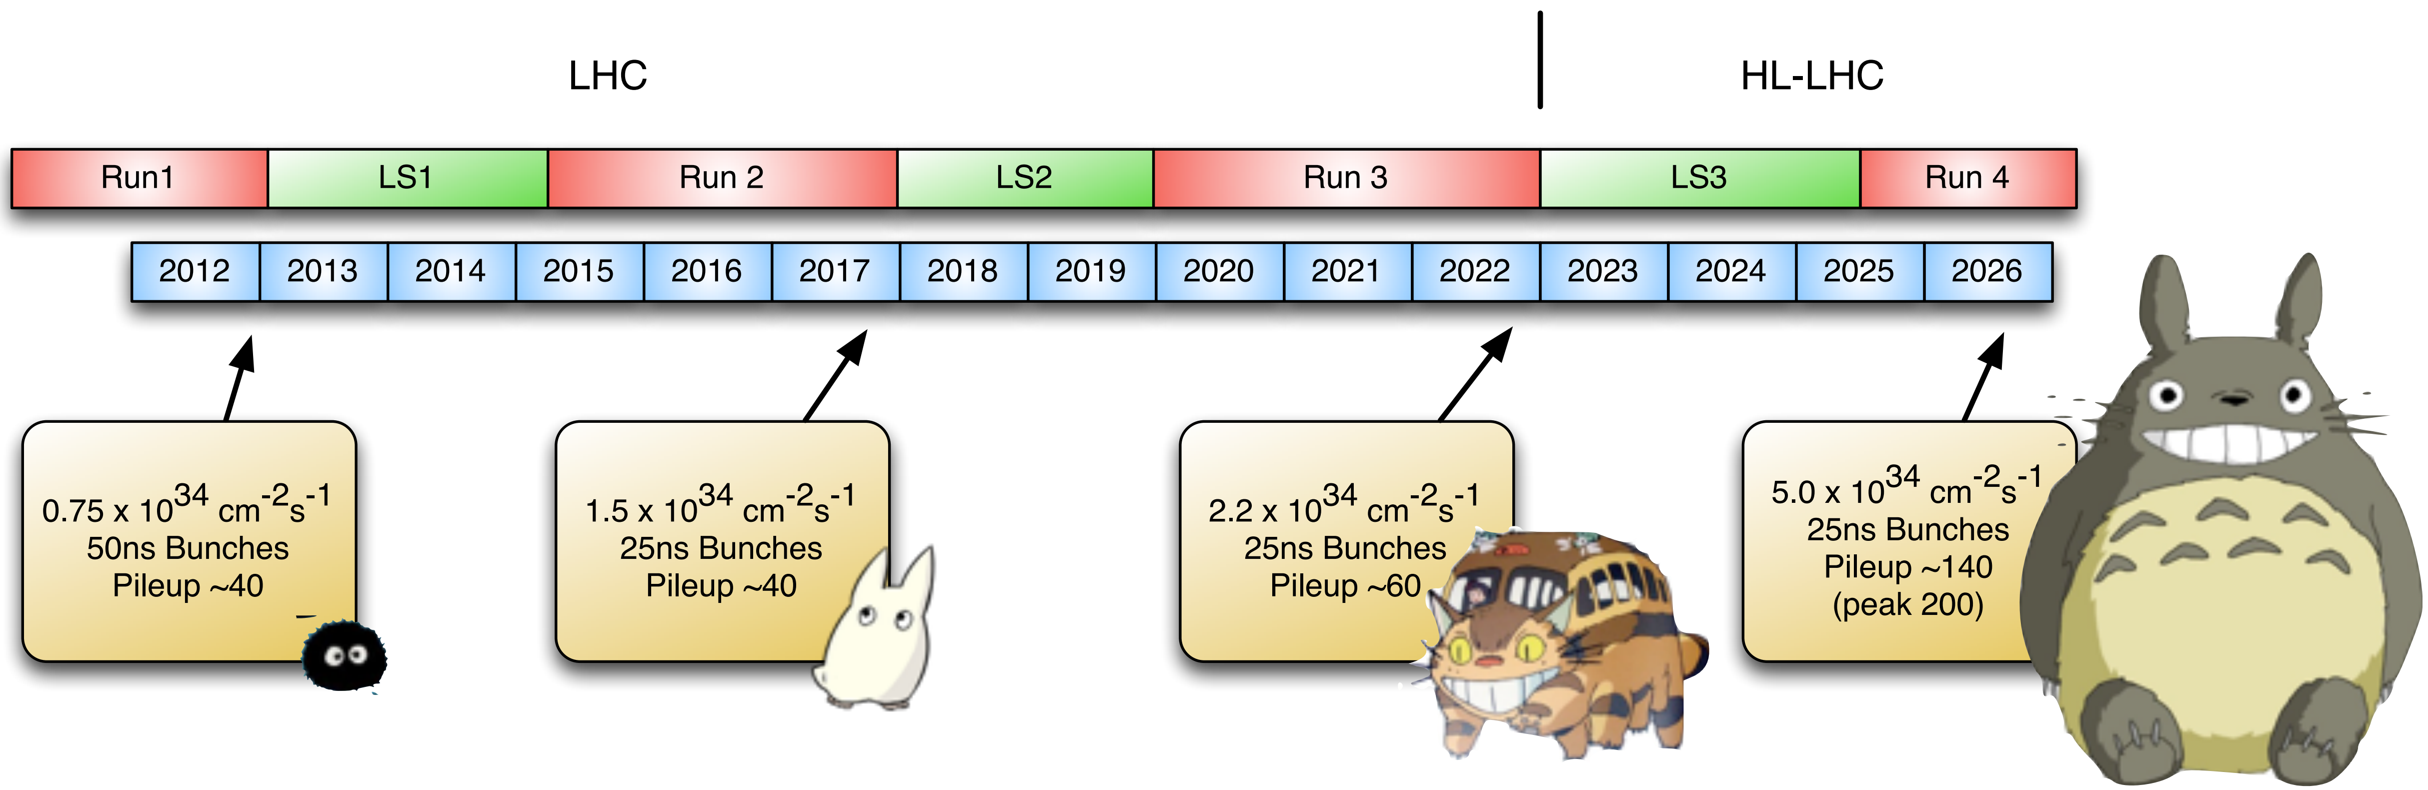

- Steady increase in machine luminosity both within runs and between runs
- Ultimate goal of 3000fb^{-1} in 10 years of HL-LHC running
  - ☞pp Collision rate of 5.6GHz
- Pileup is the most important metric of event complexity for reconstruction software

	Integrated Lumi (fb	Pileup for GPDs
Run 1	25	25
Run 2	100	40
Run 3	300	60
HL-LHC	+300 per year	140

Increasing Complexity of Events at EF Exponentially Increasing Resource Needs

- each recorded interaction consists of a hypothesis-dependent complex hierarchy of data structures
  - 2d-hits vs 3d hits vs track elements vs particle hypothesis etc
  - raw and reconstructed signals etc

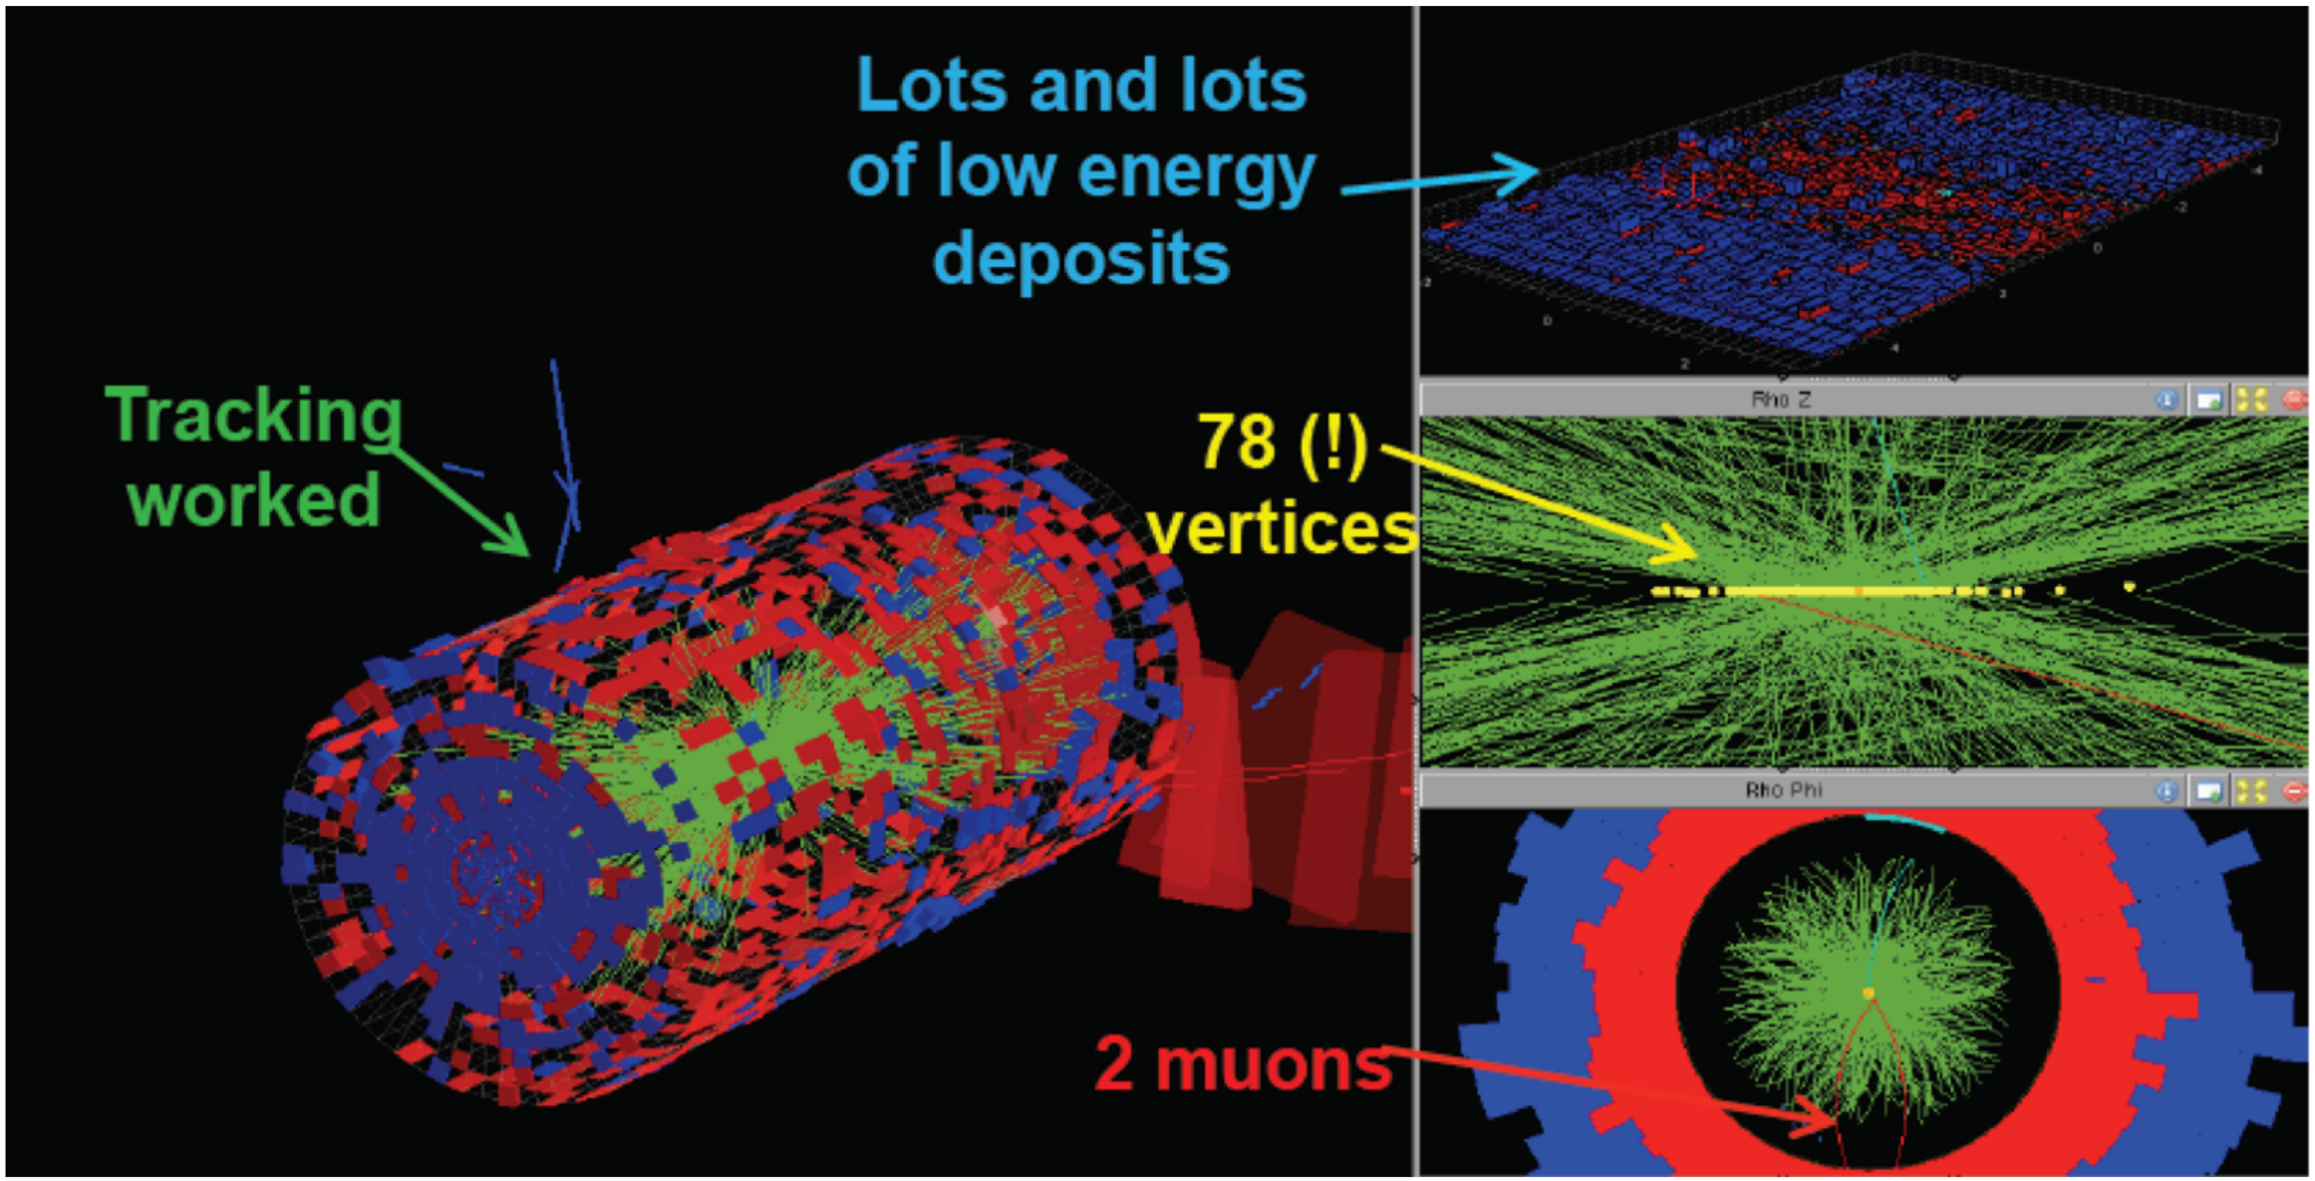

The High-Luminosity LHC will bring new Computing Challenges to the Energy Frontier

- The HL-LHC program will likely have a 10-fold increase in trigger rate and data complexity
  - ✦ computing costs will constrain data rate, and thus the possible choices on triggers and analyses
- LHC raw data: ~15 PB now; ~130 PB in 2021
  - ✦ the total dataset sizes could be up to **10x larger**, once we include processing steps and simulated data
  - ✦ data management must become much more efficient
- With flat budgets (optimistic) we might gain a factor of 4-5 in capacity over the coming decade
  - in the past LHC computing world-wide added ~25k processor cores and ~34 PB of disk, each year
- In future need to make better use of resources as the technology evolves
  - we're still not optimized, and don't use the full capacity
  - Adapting to new processor architectures becomes more challenging, requires specialized (and valuable) expertise
- Storage is cost driver, disks get cheaper only slowly
  - ✦ not all data needs to be on disk — $10 puts 1M additional events on tape (CMS) — advanced data caching, workflows, data access etc: Big Data technologies!

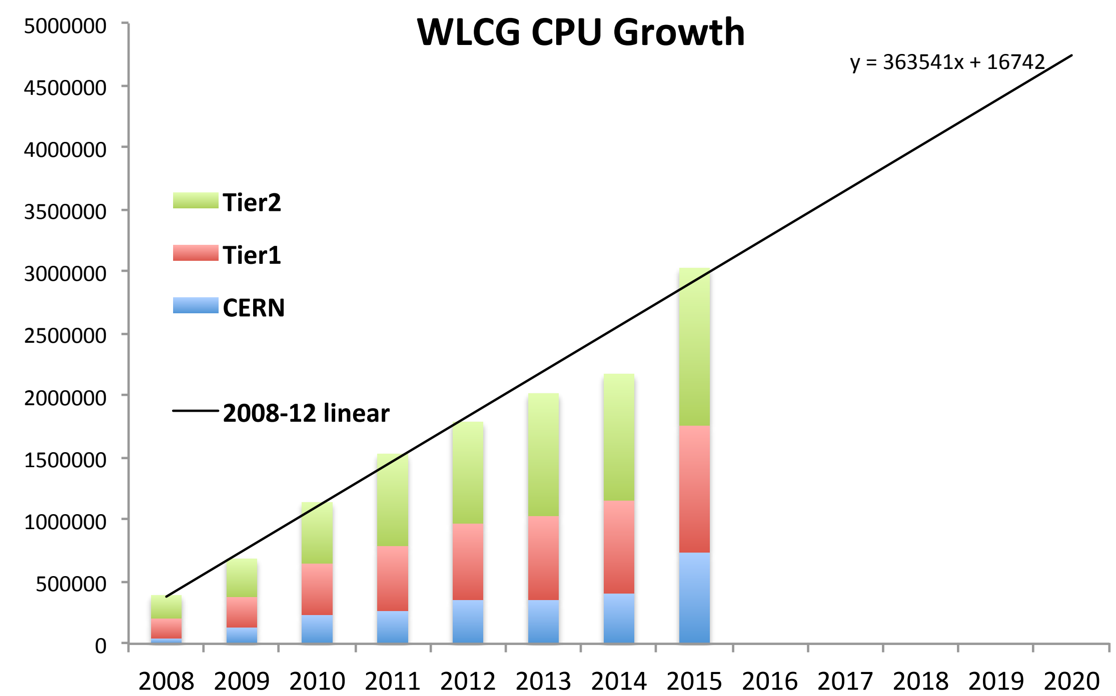

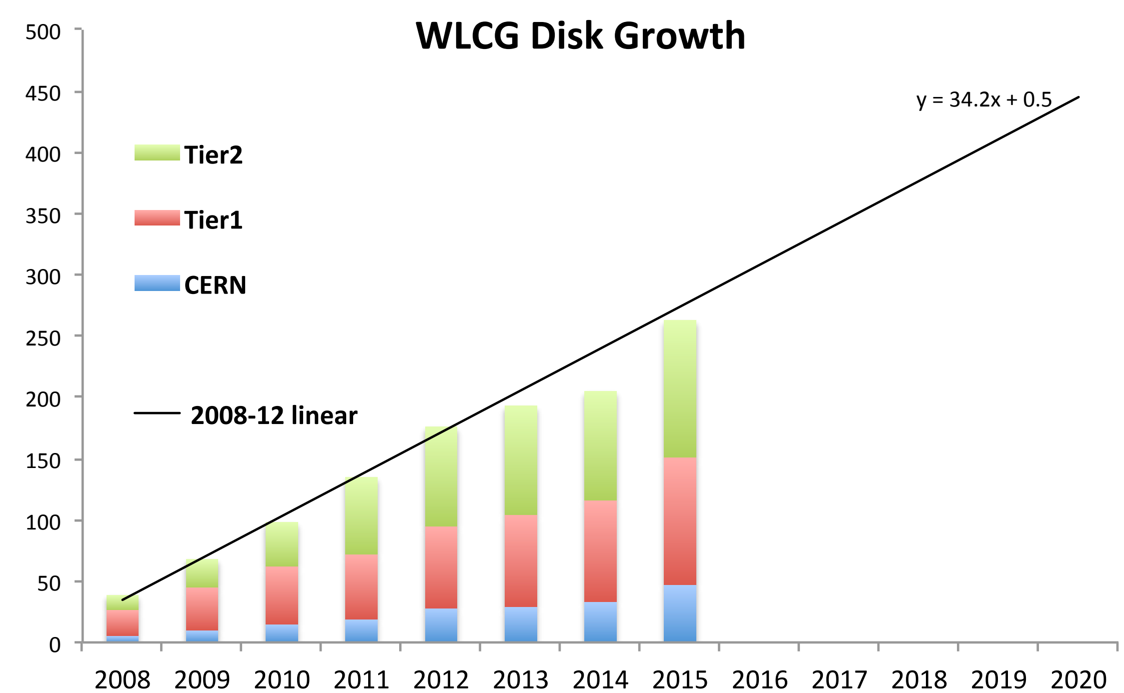

Transatlantic Networking (TAN)

- ❖ In 2012 USLHC OPMs charged a Project Execution Team (PET) with membership from US ATLAS and US CMS to recommend continuation of TAN service
- ❖ PET gathered usage information, compiled requirements, predicted needs Run1 → 2
- ❖ After a thorough evaluation process based on 3 vendor proposals, the PET has provided a detailed written recommendation to US LHC Operation Program Managers to implement the ESnet proposal (Implementation Details on slides 39 – 41)
- ❖ US LHC Operation Program Managers (and OHEP Program Managers) endorsed the recommendation to go with the proposal submitted by ESnet

Status

- ❖ ESnet transatlantic network infrastructure is completely deployed and in production
- ❖ LHCOPN transition for US ATLAS Tier-1 and US CMS Tier-1 completed on 12/12/2014
  - CERN-US Tier-1 Optical Private Network (OPN) traffic moved to ESnet infrastructure
  - ESnet/USLHCNet jointly agreed on turndown of USLHCNet services:
    - USLHCNet transatlantic circuits decommissioned in January
- ❖ Work in progress on expanding ESnet's LHCONE service to select US Tier-2 and Tier-3 sites

TAN Service Levels

- The following table summarizes current service levels under various failure scenarios. As fixing submarine cable failures may take weeks LHC traffic is distributed across 4 circuits operated on 4 independent cable systems. The table shows the impact of 1, 2 and 3 circuit/cable system outages.

	All Circuits Up		1 Circuit Down		2 Circuits Down		3 Circuits Down	
Service	**Reserved BW**	**Peak BW**	**Reserved BW**	**Peak BW**	**Reserved BW**	**Peak BW**	**Reserved BW**	**Peak BW**
LHCOPN BNL	40	200	20	100	20	40	10	40
LHCOPN FERMI	40	200	20	100	20	40	10	40
LHCONE	120	200	60	100	50	100	10	40
ESnet IP	100	100	100	100	50	100	10	40
A: 40G BOST-AMS **B: 100G NEWY-LOND** **C: 100G AOFA-LOND** **D: 100G WASH-CERN**	A: Other B: LHCOPN/LHCONE C: ESnet IP D: LHCOPN/LHCONE		Similar of B or D fail A: Other B: Down C: ESnet IP D: LHCOPN/LHCONE		Both B&D Fail A: LHCOPN B: Down C: ESnet IP & LHCONE D: Down		B, C & D Fail A: All Services B: Down C: Down D: Down	

Findings from the "Snowmass Computing Study" For the Energy Frontier and the Large Hadron Collider

- Computing resource limitations already reduce the amount of physics data that can be analyzed
- The planned upgrades (HL-LHC) are expected to result in a
  - ten-fold increase in the number of events
  - and a ten-fold increase in event complexity.
  - LHC produces about 15 petabytes (PB) of raw data per year now, but in 2021 the rate may rise to 130 PB
- Efforts to increase code efficiency, parallelism, data processing
  - explore the potential of computational accelerators
  - advance from sequential to "big-data" type data analysis
- More than half of the computing cost is now for storage
  - in future it may be cost-effective to recalculate, rather than store
- Attention on data management and wide-area networking
  - assure network connectivity for distributed event analysis

Computational Challenges for the LHC

- Challenging resource needs require efficient and **flexible use of all resources**
  - We're proactively looking into ways of tapping into new kinds of resources
  - both **Distributed High-Throughput Computing** (Grids, Clouds)
  - and **High-Performance Computing** (some successes with NERSC, LCF, XSEDE)
  - **Sharing** and **opportunistic use** help address resource needs, from all tiers of computing, and now including community or commercial clouds etc
- To stay on the Moore's law curve, need to proactively make full/better use of **advanced architectures:** multi-threading, GPU environments, low-energy CPUs
  - With the need for more parallelization the **complexity of software and systems** continues to **increase**: frameworks, workload management, physics code
  - Important needs for **developing and maintaining expertise** across offline, computing, POGs etc, including re-engineering of frameworks, libraries and physics codes, adapting key software tools
- Unless corrective action is taken we could be **frozen out of cost effective computing solutions** on a time scale of 10 years.
  - There is a large code base to re-engineer
  - We currently do not have enough people trained to do it

R&D and Improvements Needed

- In software, HEP has made good progress on new architectures
  - low cost and low power ARM processors, high performance GPU and co-processor systems (e.g. GPU-based photon tracking in IceCube,
  - need to deal with much lower memory/core
- Substantial development effort to improve ability to run across many cores
  - thread safe code and libraries, algorithm re-engineering
  - matching the multi-core architecture, but running against Amdahl
- New resources and more techniques in resource provisioning
  - developing access to opportunistic computing and migrating to cloud provisioning tools
- More efficient use of storage
  - move to SSD, dynamic data placement, more reliance on data served remotely and content delivery networks
  - subsequently an even larger reliance on networks (see M.Ernst)
- These new developments will allow more flexible offerings for more computing and storage capacity at lower cost and energy

from M.Ernst

Cloud Provisioning

Example: AutoPyFactory

The deployed hardware is expected to scale to 100k concurrent jobs

- Experience with ~30k jobs
- Setup serves serial and multi-core queues
- Policy-based VM lifecycle management

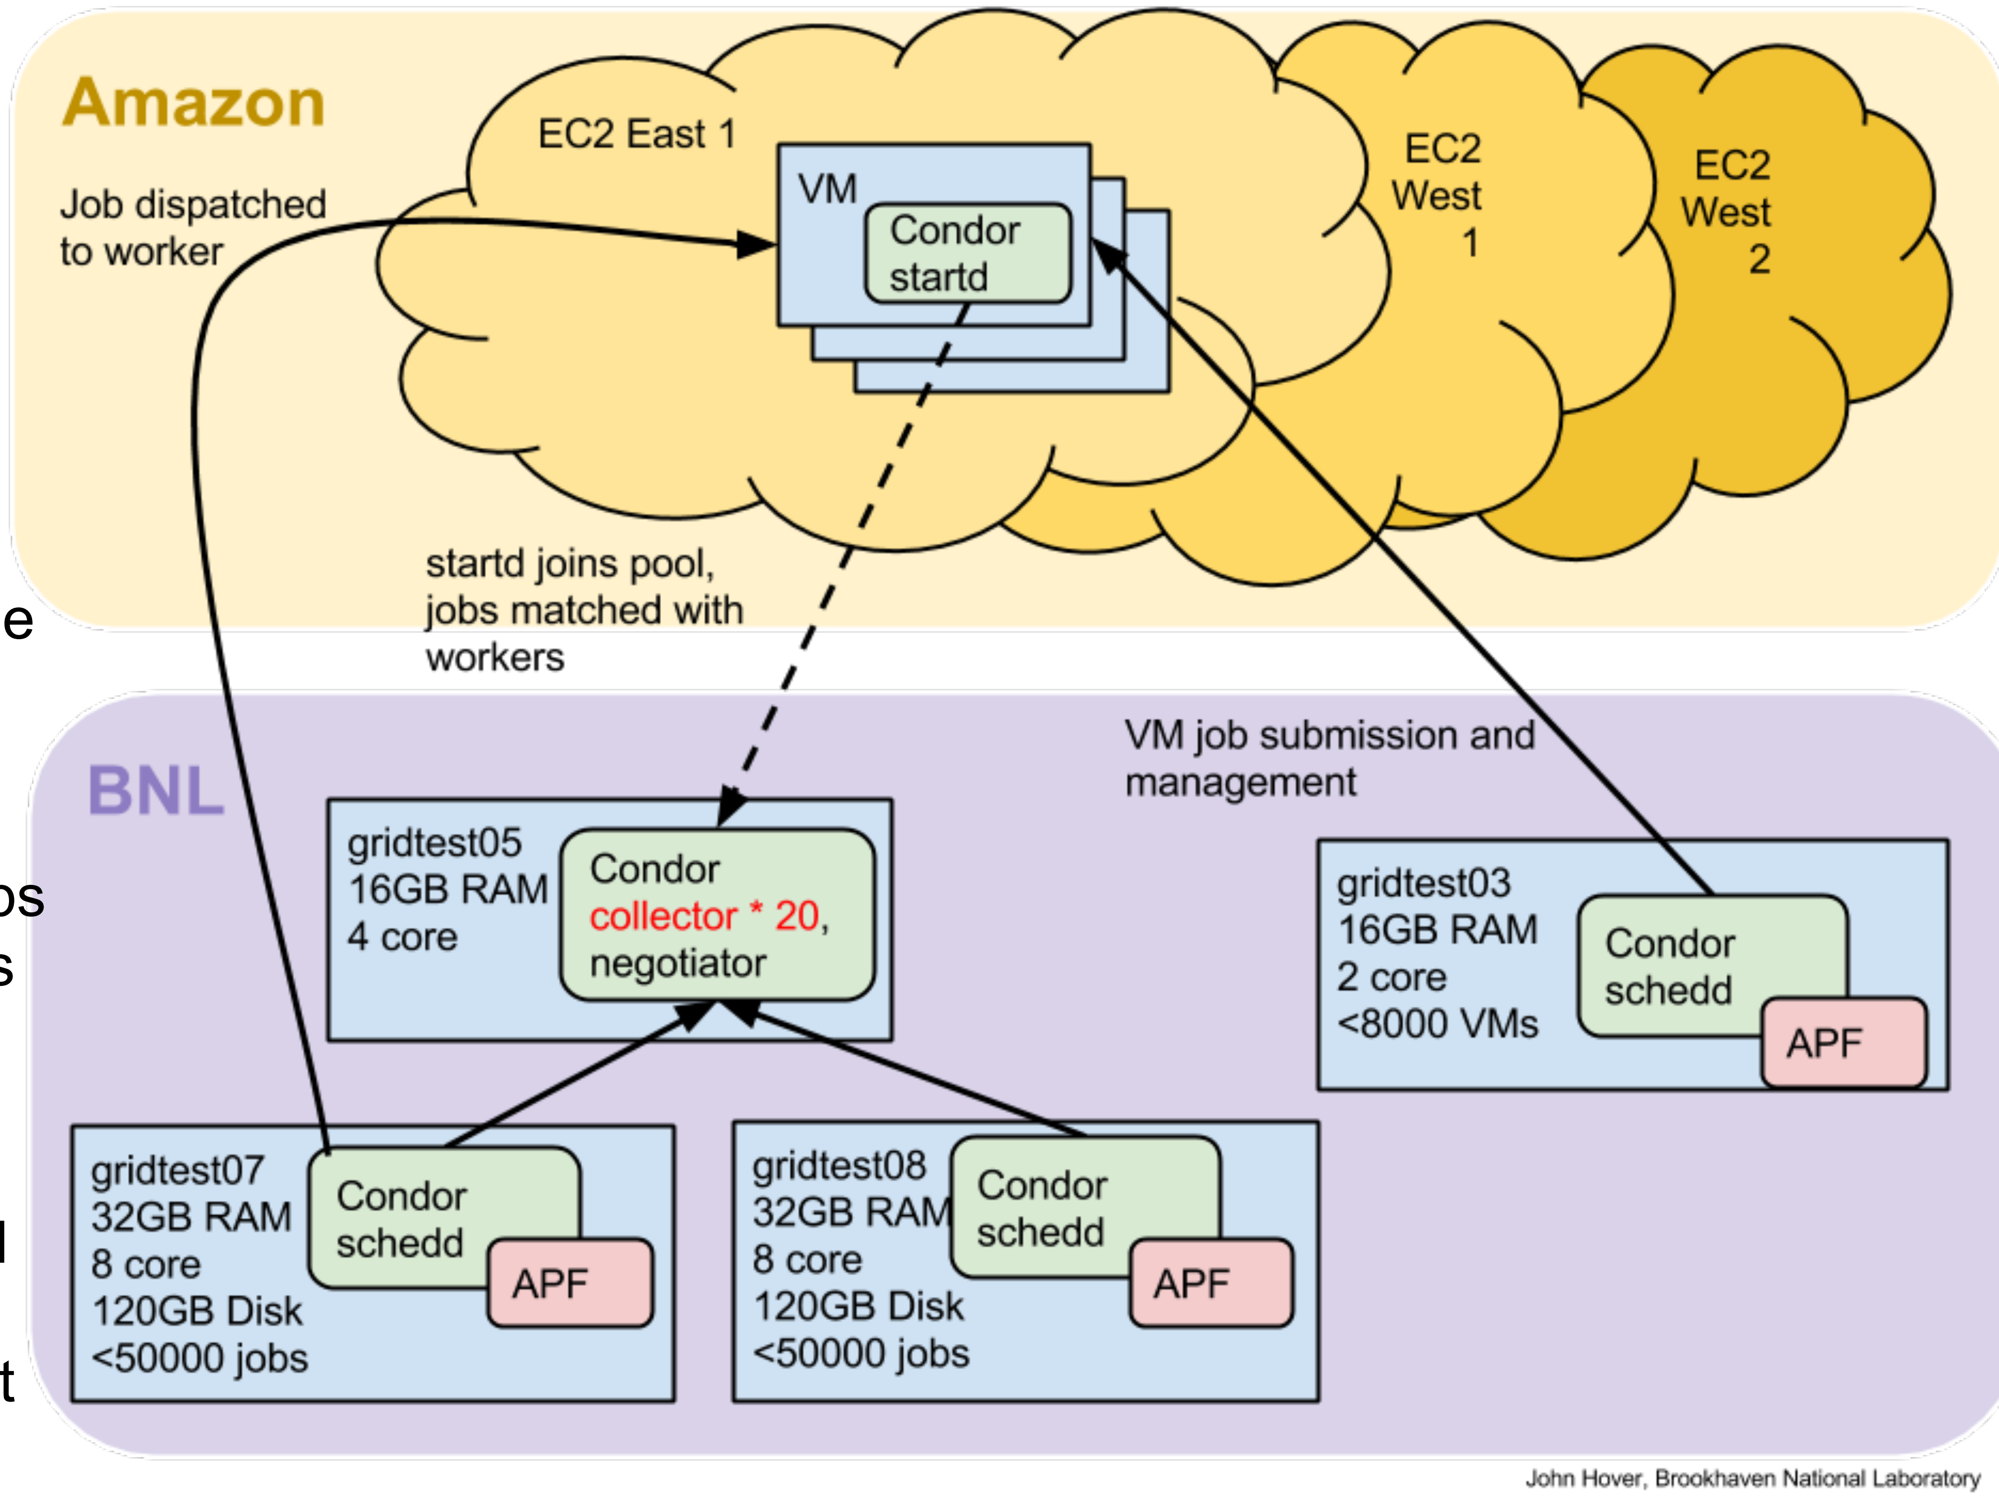

from M.Ernst

Utilize Cloud for Compute

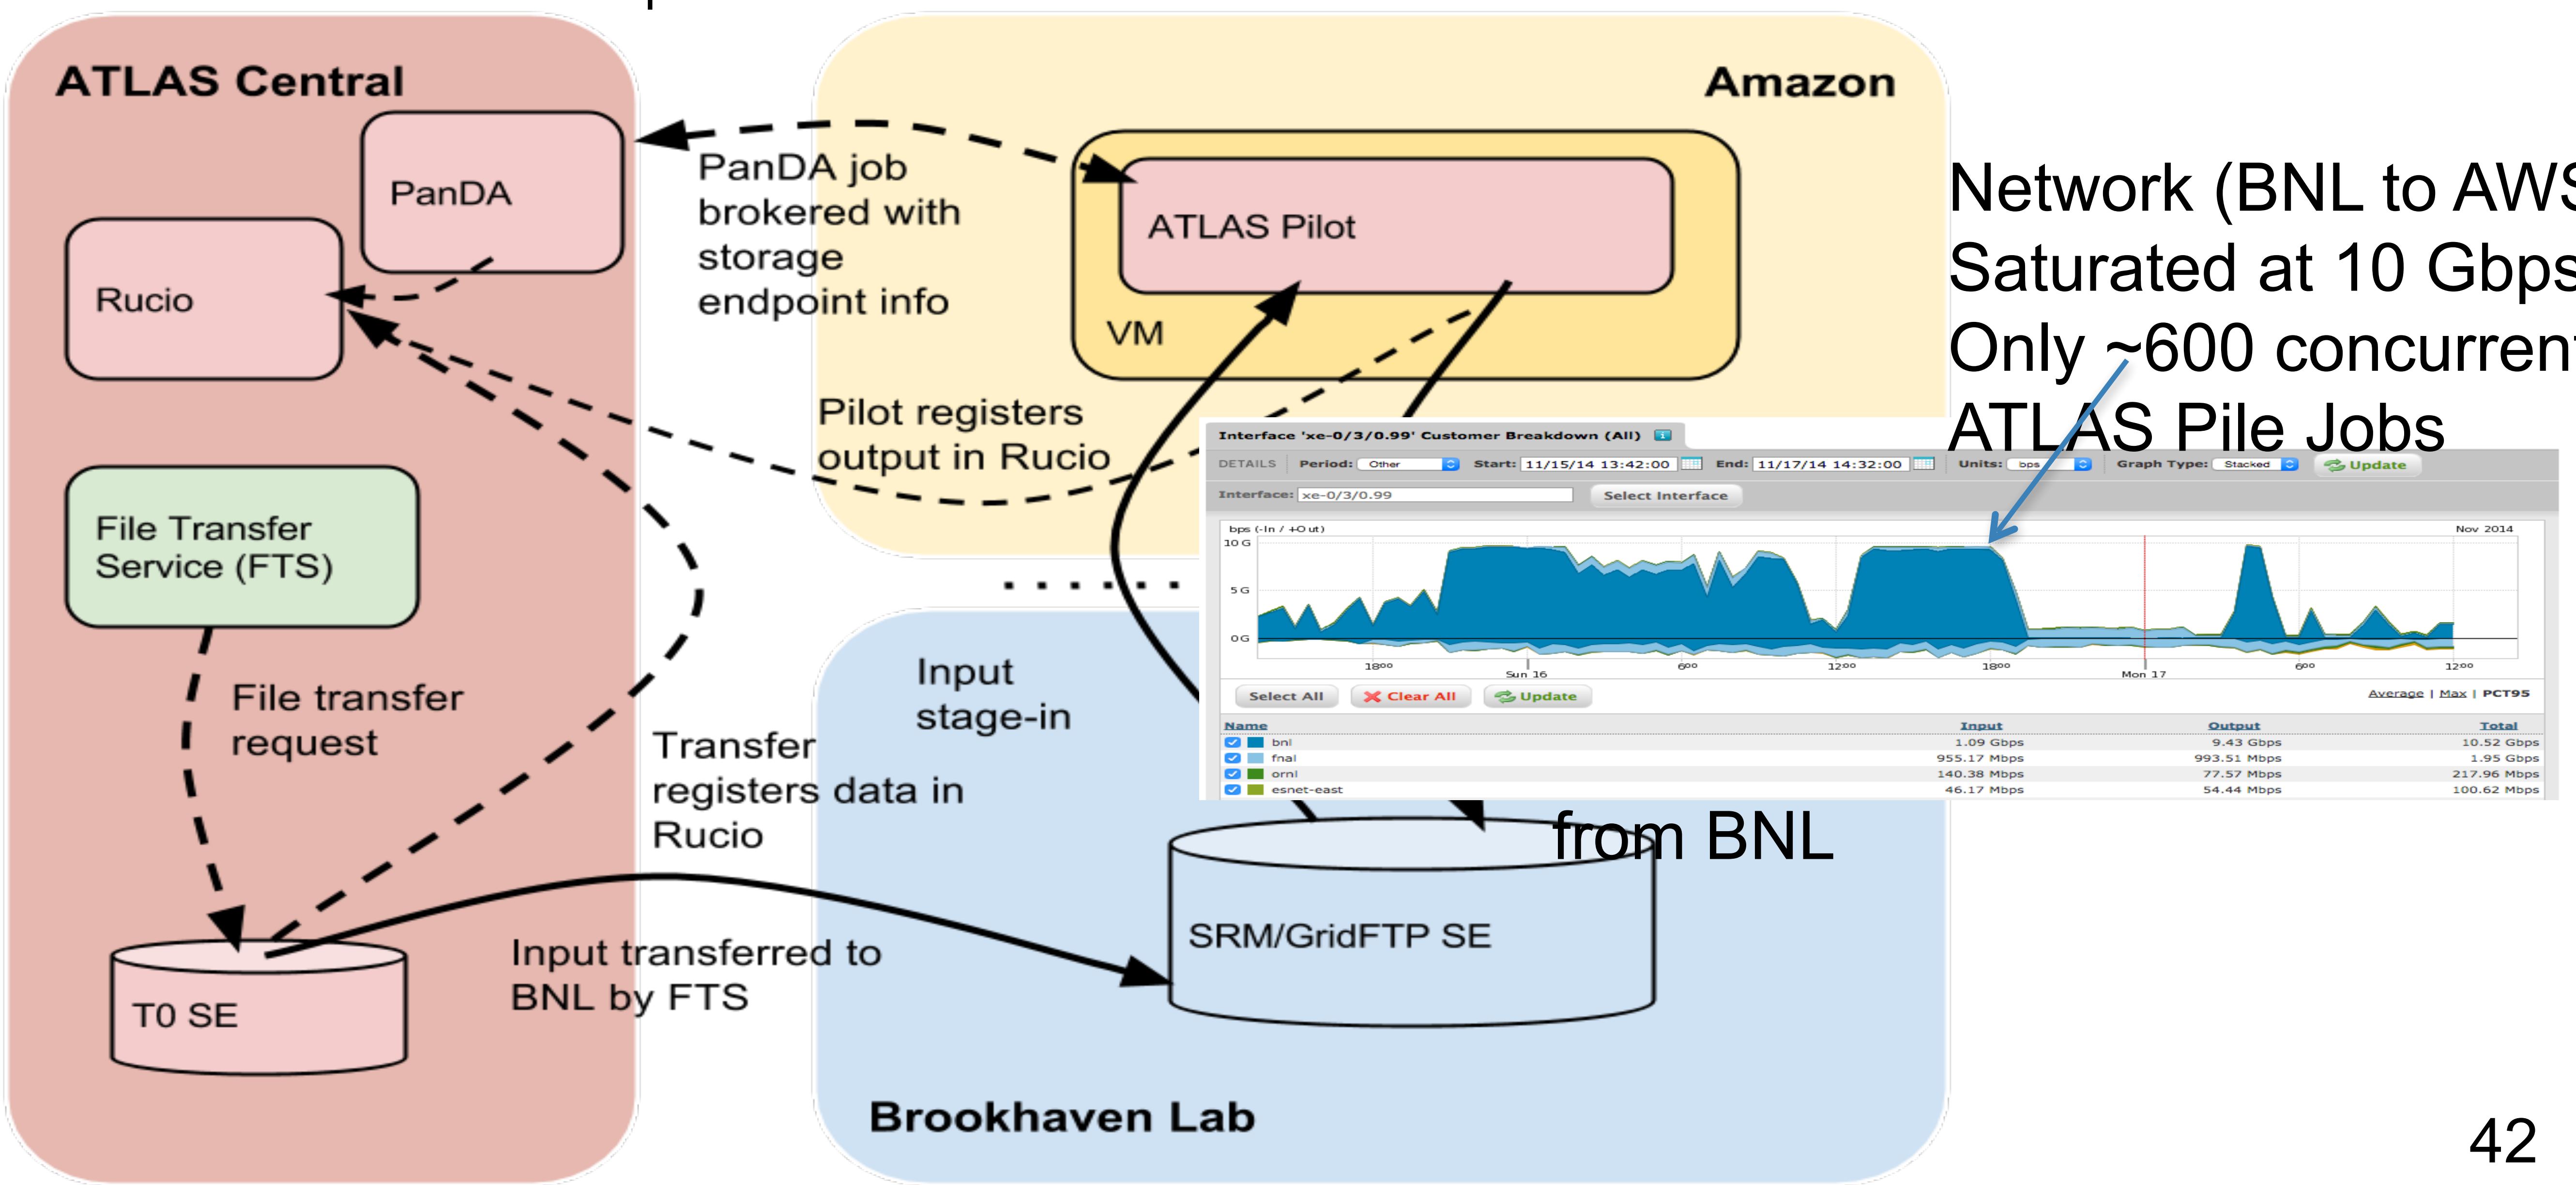

from M.Ernst

Using the Cloud for Compute and Storage

Long-term goal: Run entirely within EC2, with link limitations only affecting SE to SE transfers. “Site” stage-in/out via S3.

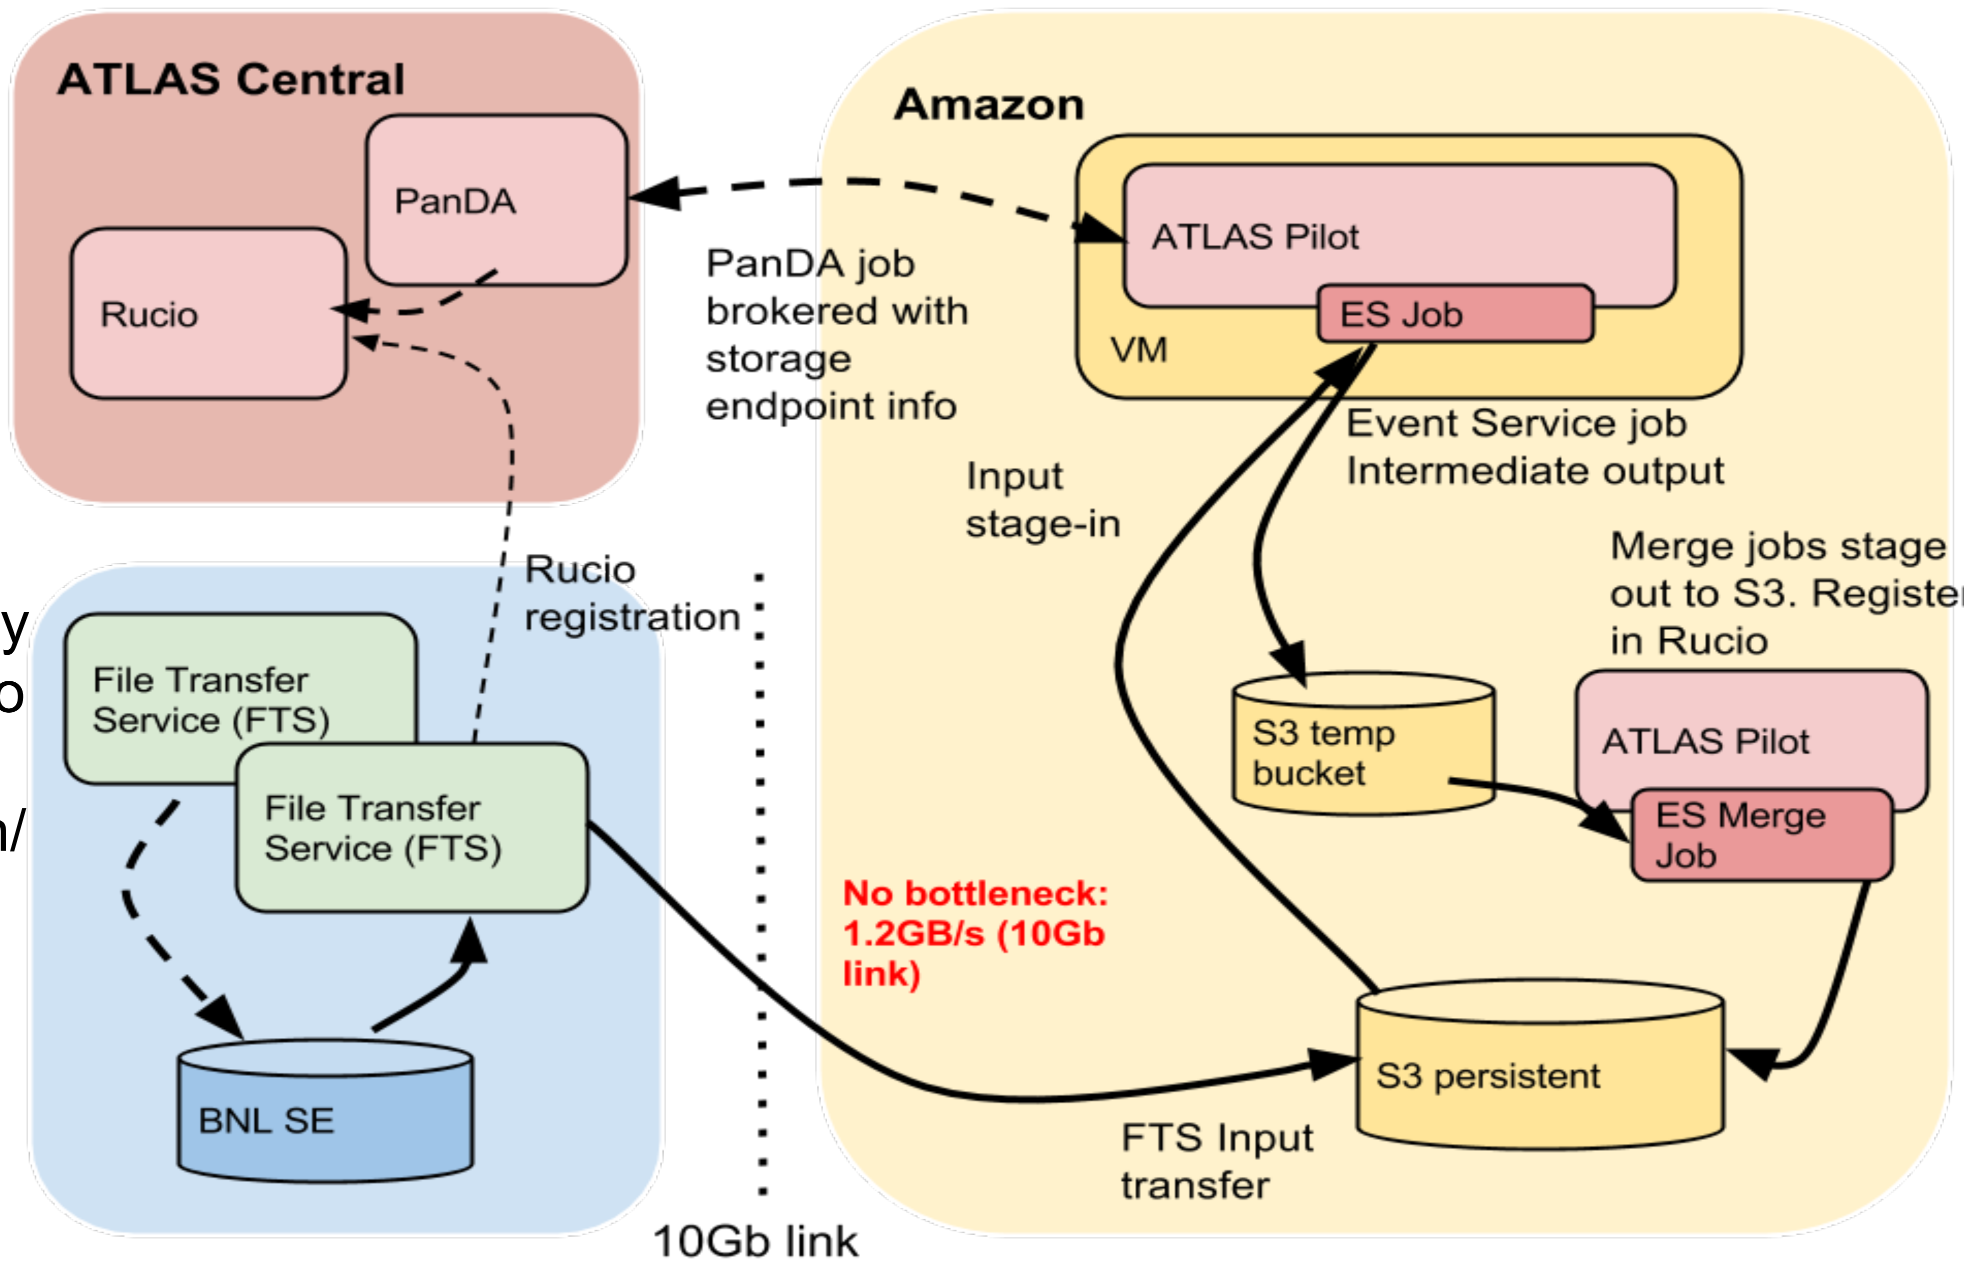

Jobs only using S3, no scaling issues.

The Big Data Frontier   from Wired Magazine

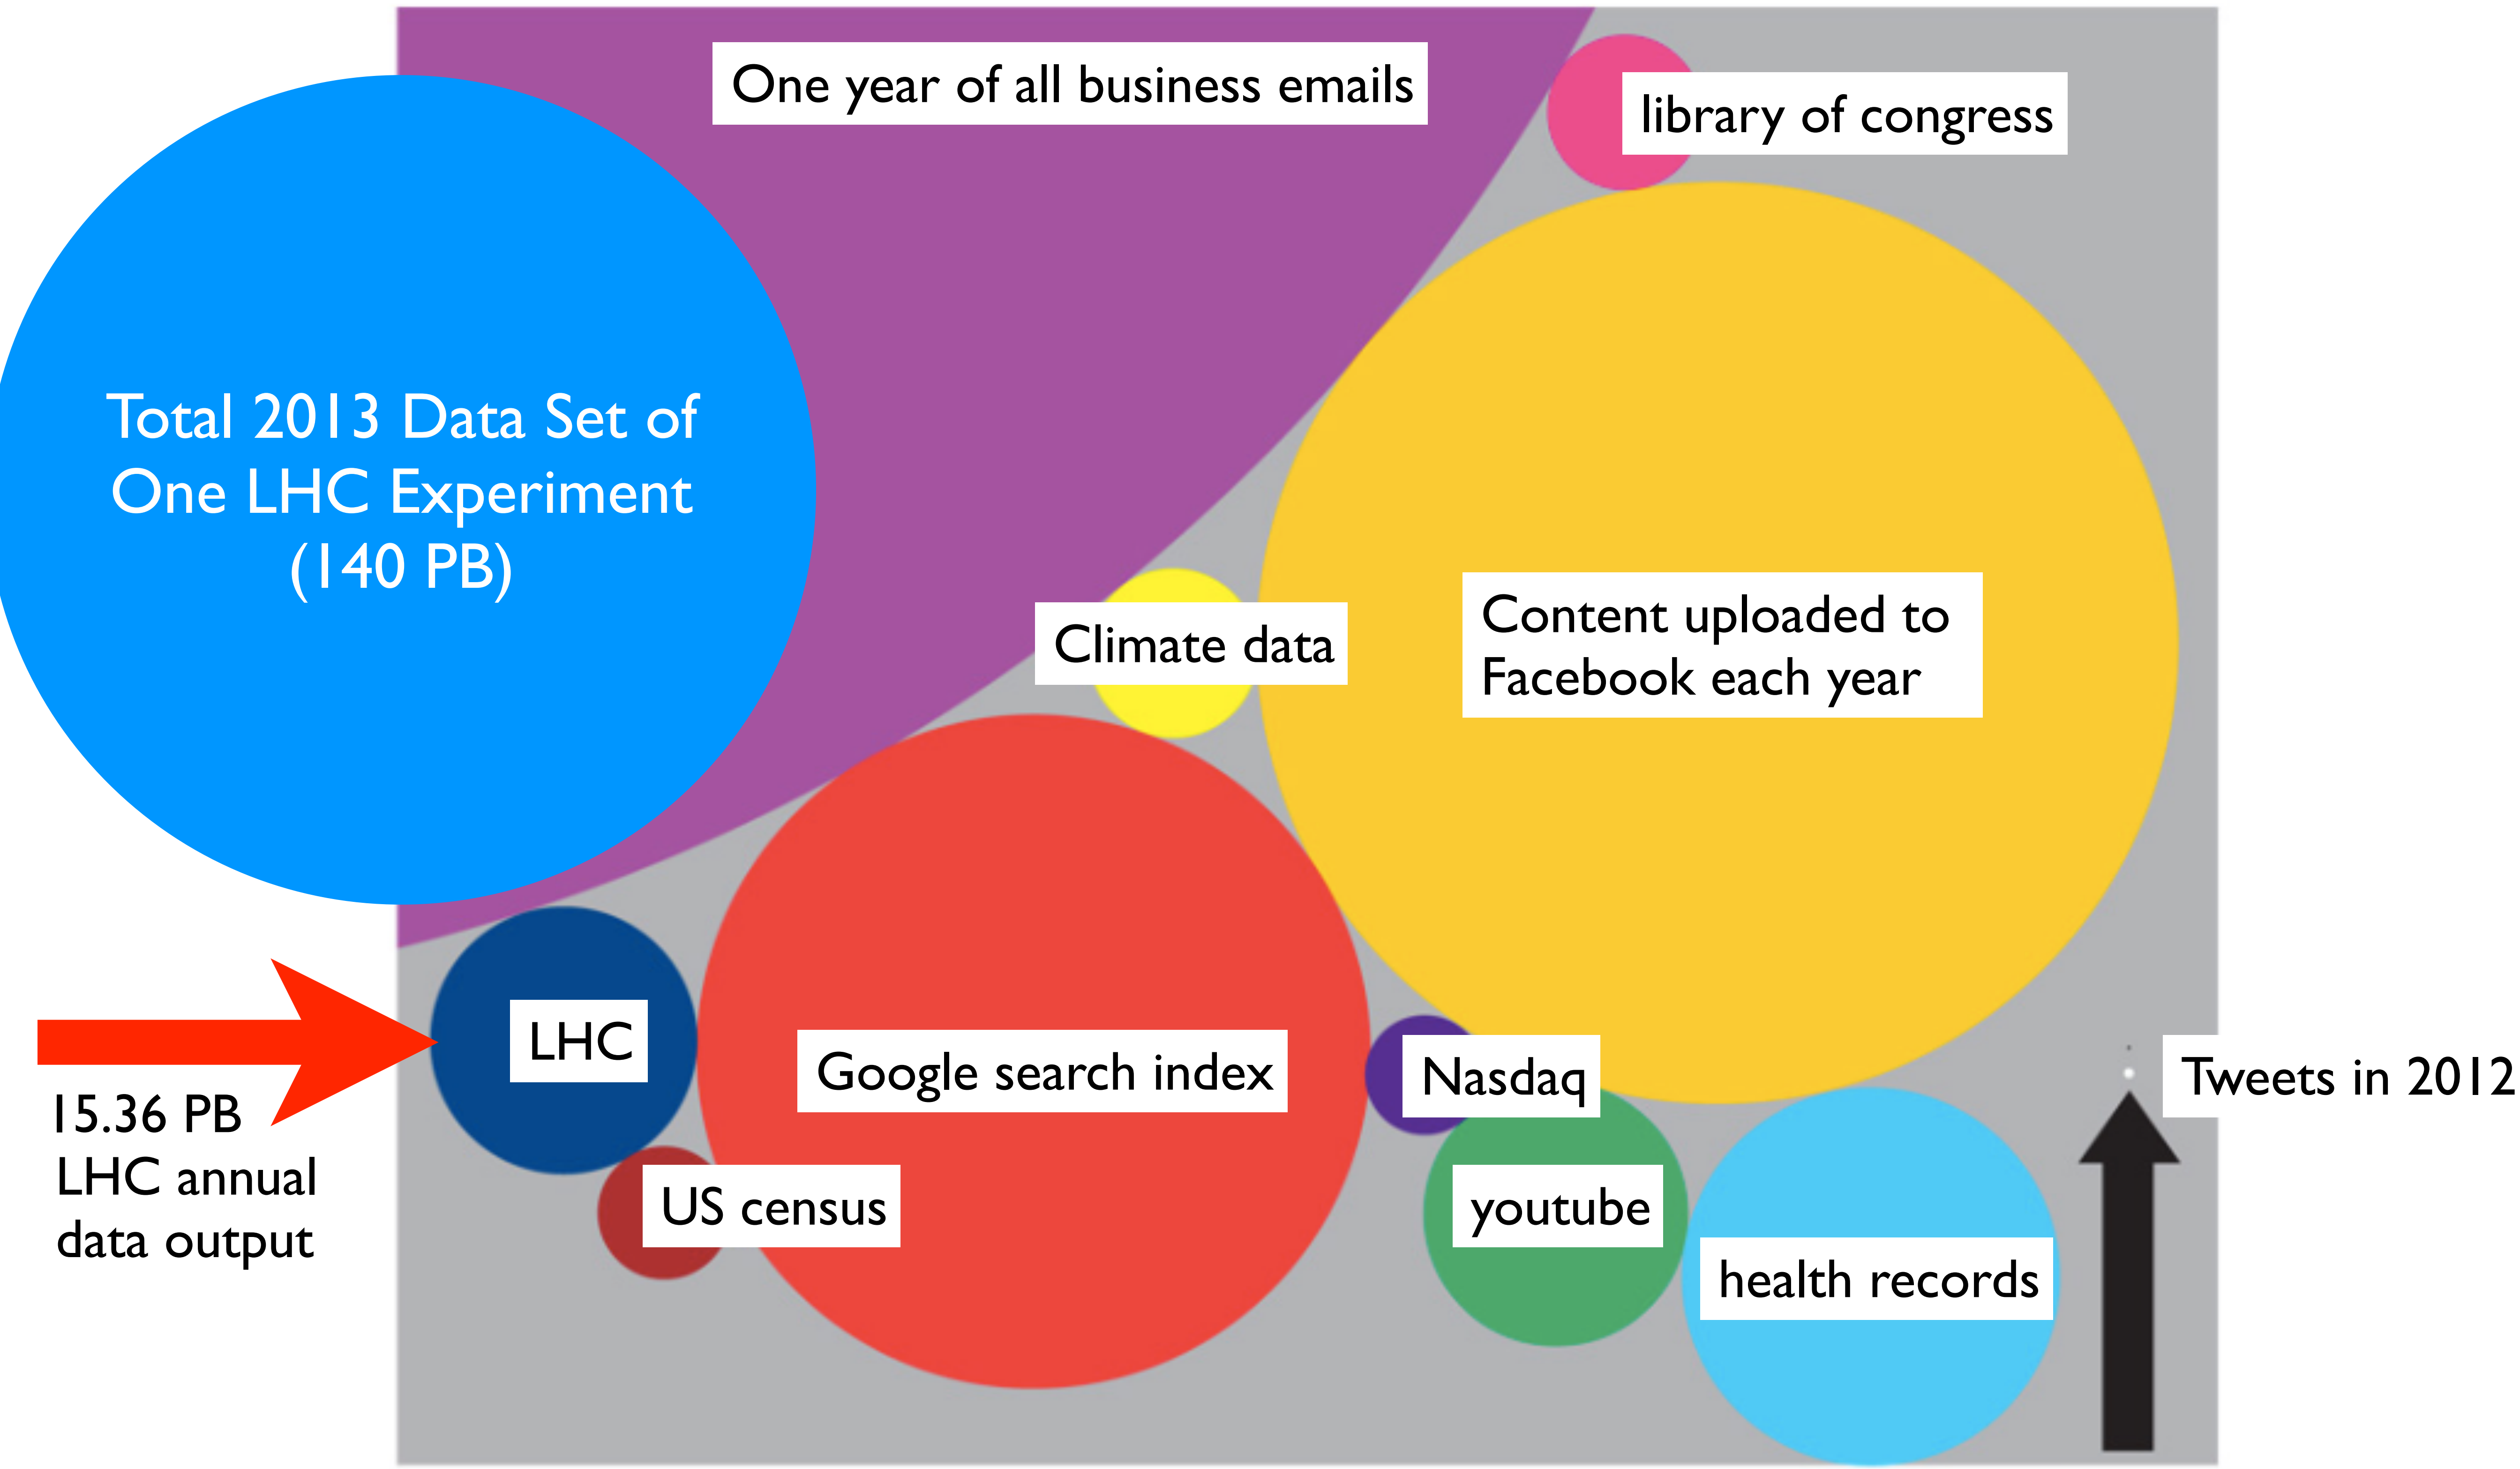

http://www.wired.com/magazine/2013/04/bigdata/